

UNIVERSITEIT VAN AMSTERDAM



MSc Physics and Astronomy (joint degree)
Advanced Matter and Energy Physics track

Master Thesis

Towards Ultracold Strontium Atom Transport for a Continuous Superradiant Laser on the mHz Transition

Submitted by

Elisa Wenzel

Daily Supervisor: Benedikt Heizenreder
Supervisor: prof. dr. Florian Schreck
Second Examiner: dr. Rene Gerritsma

September 2025 – July 2026

Van der Waals-Zeeman Institute
University of Amsterdam

Abstract

Continuous superradiant lasing on the mHz-linewidth clock transition in strontium requires a source of atoms that is simultaneously cold, dense, and continuously replenished with a stable flux. This thesis presents contributions to the design, setup, and characterization of key stages of the atom preparation and transport needed to realize such a source.

A dual-position narrow-line magneto-optical trap (MOT) was optimized and characterized. It provides a continuous source of atoms with a loading rate of 1.5×10^8 atoms/s and a saturation atom number of 3.7×10^8 in the broadband stage at temperatures of 8.0 µK. The spatially separated single-frequency stage reaches lower temperatures, here 3.3 µK. A high-power optical dipole guide can be used to transport the atoms from the single-frequency MOT to a separate science chamber. The required powers were determined theoretically and the experimental implementation of the guide was completed. In the science chamber, the atoms will be slowed down by molasses beams and recaptured in a reservoir trap, both of which were set up as part of this thesis. Finally, an active laser intensity stabilization system was tested and characterized, and its role in the experiment discussed. The best-performing configurations suppress the relative intensity noise by over 25 dB at low frequencies.

CONTENTS

List of Optical Elements **3**

1 Introduction **4**

2 Laser Cooling and Optical Potentials **7**

2.1 Light Forces on Atoms and Laser Cooling 7

2.1.1 The Scattering Force 8

2.1.2 The Dipole Force 10

2.2 Gaussian Beams and Beam Shaping 11

2.3 Intensity Noise and Heating 11

3 Experimental Setup for Continuous Superradiance on Ultracold Atoms **13**

4 Towards Continuous Transport of Ultracold Strontium **17**

4.1 Blue and Red MOT Cooling to µK Temperatures 17

4.2 Optical Dipole Guide and Atomic Transport 22

4.3 Slowing and Recapture in Reservoir Trap 28

4.4 Supporting Optical Systems 32

4.4.1 689 nm Light Distribution and Frequency Shifting 32

4.4.2 Retroreflection Beam Setups 33

5 Laser Intensity Stabilization **36**

5.1 Proportional-Integral-Derivative Controller and Experimental Setup 37

5.2 Step Response Analysis 40

5.3 Noise Spectra Analysis 42

5.4 Sample and Hold Function 46

6 Conclusion and Outlook **49**

LIST OF OPTICAL ELEMENTS

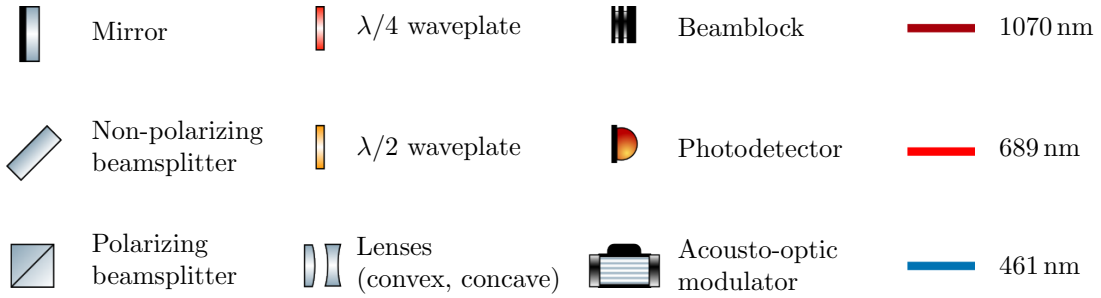


Figure 1: List of optical elements and wavelength colors used throughout this thesis. Symbols from the ComponentLibrary by Alexander Franzen [1].

1 | INTRODUCTION

Frequency is the physical quantity that can be measured with the highest precision, and time-keeping directly builds on this. The second itself is defined by the frequency of a specific microwave transition in the cesium atom [2]. Optical transitions oscillate several orders of magnitude faster than this microwave reference, which increases the achievable resolution provided the transition remains equally narrow [3]. This is why optical clocks are expected to replace the cesium standard in a future redefinition of the second [4]. Atoms are natural candidates for a stable reference, since the frequencies of their internal transitions arise fundamentally from quantum physics rather than from properties of a fabricated system. The transition frequency of an atom, however, only defines a reference and it does not by itself provide a macroscopic oscillating signal that can be counted. A clock therefore additionally needs an external oscillator that provides the countable signal and is sufficiently stable itself. State-of-the-art optical clocks realize the reference either as a single trapped ion or as an ensemble of neutral atoms confined in an optical lattice [5, 6], while the oscillator is a laser stabilized to an ultra-stable cavity. This architecture has two important limitations. The first limitation is the imperfect stability of this oscillator. The reference cavity is mechanically and thermally isolated, but residual vibrations and thermal drifts of the mirrors become frequency noise on the light, and even a perfectly isolated cavity remains limited by Brownian thermal noise [7]. The second limitation is the cyclic operation of these clocks, where an atomic ensemble must be prepared new for every measurement, leaving dead time in between.

Superradiance on a narrow optical transition offers an alternative approach [8]. Dicke predicted in 1954 that an ensemble of atoms decaying collectively rather than independently emits with a spontaneous emission rate enhanced by the number of atoms [9]. A cavity can induce this collective behavior by mediating the atom interactions through a common mode. How much this cavity influences the output frequency depends on the regime of operation. A conventional laser system operates in the good-cavity regime, where the cavity linewidth is narrower than the atomic linewidth and the cavity dictates the output frequency. In the bad-cavity regime this is reversed and the cavity linewidth greatly exceeds that of the atomic transition, so the output frequency is set by the atoms and cavity-length fluctuations are suppressed [10], as illustrated in Figure 2. The frequency reference is then given by the atoms themselves, analogous to a hydrogen maser, but on an optical rather than a microwave transition [11]. In this regime the emission linewidth can even fall below the natural linewidth of the atomic transition [12], as demonstrated for ^{88}Sr [13], giving a further advantage for precision measurements.

Most precision measurements with cold atoms are performed in cycles in which an ensemble is cooled, probed, and destroyed before a new ensemble is prepared for the next repetition. No measurement can be taken during the preparation stage of each cycle. During this dead time, the frequency fluctuations of the local oscillator are not measured and therefore cannot

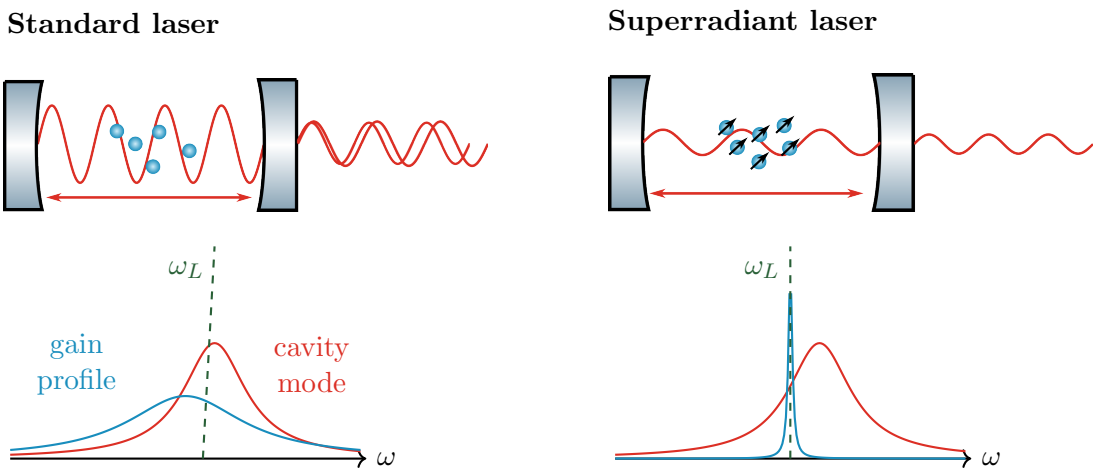


Figure 2: Comparison between a conventional laser operating in the good-cavity regime (left) and a superradiant laser operating in the bad-cavity regime (right). Top: atoms confined in an optical cavity. In the superradiant case, the atomic dipoles synchronize via the cavity-mediated interaction. Bottom: corresponding linewidths of the atomic gain profile (blue) and the cavity mode (red), with the laser output frequency ω_L indicated. In the good-cavity regime the narrower cavity mode sets the output frequency, making it sensitive to cavity-length fluctuations. In the bad-cavity regime it is pulled towards the narrower atomic line instead, strongly suppressing this sensitivity.

be corrected, so they alias into the clock output. This is known as the Dick effect, and it limits the achievable stability of an atomic clock even when each individual measurement is highly precise [14]. In cycled operation, the instability then averages down as $1/\sqrt{\tau}$ with averaging time τ. A continuous active clock output removes this dead time, allowing the local oscillator to be monitored without interruption and effectively phase-locked to the atoms. Uninterrupted phase tracking has been predicted to enable stability scaling as fast as $1/\tau$ [15].

Combining these advantages gives an intrinsically stable and continuously operating frequency reference. This is the goal of the experiment this thesis contributes to, a superradiant laser in the bad-cavity regime, in which some atoms are always being prepared while others are always emitting, so that no dead time is introduced.

The alkaline-earth metal strontium has a level structure well suited to this plan. It is a standard choice in ultracold experiments, with transitions convenient for laser cooling and trapping. Importantly for this work, it offers two narrow transitions suitable for bad-cavity operation: the spin-forbidden ${}^1S_0 \rightarrow {}^3P_1$ transition, with a linewidth of a few kHz, and the doubly forbidden ${}^1S_0 \rightarrow {}^3P_0$ clock transition, with a linewidth at the mHz level. For metrological applications, the clock transition is more interesting, because a narrower linewidth gives a more precise measurement of the atom's oscillation frequency. Pulsed superradiance has already been demonstrated on the clock transition in ^{87}Sr, with collective enhancement of the atomic decay

rate observed at the level of 10^4 [16]. What remains is making this operation continuous, which requires redesigning the experimental setup and solving the resulting physical and technical challenges.

The atom source must meet several requirements at once. The atoms must be cold enough to suppress Doppler broadening of the narrow transition, dense enough to efficiently generate superradiance in the cavity, and delivered continuously, so that the superradiant emission can be sustained. An alternative would be to keep the same atoms and repump them many times [17], but this requires managing atom losses and recooling continuously, making it difficult to keep the atom number constant. Instead, new atoms can be delivered continuously either transversely, as an atomic beam crossing the cavity axis, or longitudinally, with ultracold atoms traveling along the cavity axis. The former is simpler, using only transverse cooling of the atomic beam. However, this means the atoms have residual motion along the cavity axis that introduces a Doppler shift, restricting it to the broader kHz transition [18, 19]. The mHz transition instead requires confinement in an optical lattice and cooling to microkelvin temperatures.

Continuous preparation of ultracold strontium has been demonstrated in recent years, including a steady-state narrow-line magneto-optical trap and a continuous Bose-Einstein condensate [20, 21]. This shows that the sequential stages of a cold-atom experiment can be operated in parallel at spatially separated locations, making the second, more demanding approach feasible. The experiment in this thesis builds on the same architecture. Atoms are continuously cooled to microkelvin temperatures, guided into a separate science chamber, and recaptured in a reservoir at the cavity, from which a moving optical lattice transports them through the cavity mode.

This thesis presents parts of the design, setup, and characterization of the stages that make up this transport chain, from the continuous preparation of ultracold atoms up to their recapture at the cavity. All of these stages rely on precisely controlled laser light, where intensity fluctuations heat atoms in optical potentials and modulate the number of atoms delivered to the cavity. The final part of this thesis therefore develops and characterizes an active laser intensity stabilization system for the experiment.

Chapter 2 introduces the theoretical framework of laser cooling and optical potentials relevant to this setup. Chapter 3 gives an overview of the experimental apparatus. Chapter 4 describes and characterizes the stages of the transport chain, from the magneto-optical traps through the dipole guide to the reservoir and its supporting optical systems. Chapter 5 presents the intensity stabilization system, and Chapter 6 concludes with an outlook.

2 | LASER COOLING AND OPTICAL POTENTIALS

2.1 Light Forces on Atoms and Laser Cooling

The success of atomic physics experiments and their ability to reach quantum regimes is largely due to advances in laser cooling and trapping methods, as lasers provide a powerful and versatile tool to manipulate both the internal states and external motion of atoms. For precision experiments such as narrow-line superradiance, Doppler broadening due to thermal motion is a fundamental limiting factor, washing out transitions at the mHz level. Beyond minimizing Doppler shifts, cooling and trapping are essential to confine atoms long enough for such experiments to be conducted and control their position in the experiment.

In the following, I explain the origin of the two fundamental light forces on atoms and describe how molasses cooling and magneto-optical traps can cool and confine atoms effectively. These are two mechanisms of laser cooling that are particularly relevant for this experiment, out of many advanced techniques. The theory is developed in the semi-classical approximation, treating the electromagnetic field classically and the atom quantum mechanically. I consider a monochromatic light field with fixed polarization, and apply the electric dipole approximation, valid when the atomic size is much smaller than the optical wavelength. The atom can be modelled as a two-level system with transition frequency ω_0 and spontaneous decay rate Γ (in angular frequency units), moving at velocity $\mathbf{v}$ as shown in Figure 3(a). It is driven by a laser field with frequency ω and wavevector $\mathbf{k}$.

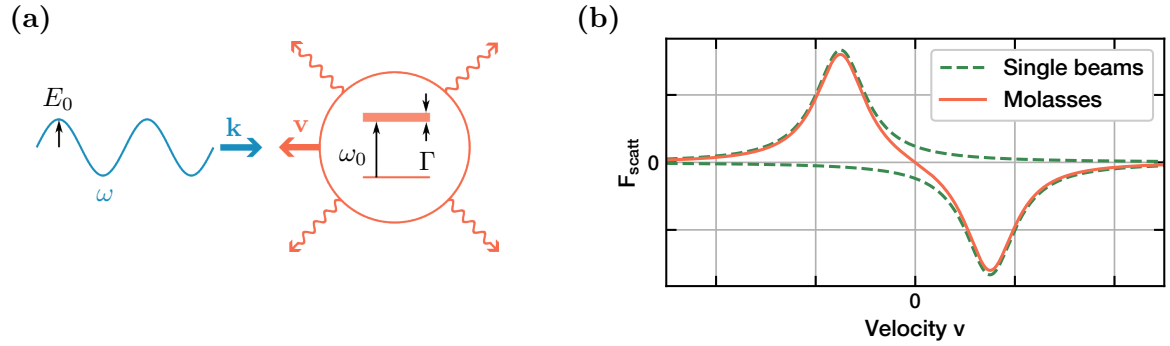


Figure 3: (a) Two-level atom interacting with a laser field, with the relevant parameters indicated. (b) Scattering force from Equation (2.1) as a function of atom velocity for each beam of a counter-propagating pair individually (green) and their combined molasses force (orange), for red detuning $\delta < 0$. The molasses force provides a velocity-dependent damping near $v = 0$.

2.1.1 The Scattering Force

The scattering force arises from the momentum transfer of scattered photons onto the atom during absorption. Its value is given by the product of the photon momentum and the scattering rate

$$\mathbf{F}_{\text{scatt}} = \hbar \mathbf{k} \, \Gamma \rho_{22}, \tag{2.1}$$

where ρ_{22} is the excited state population. The force points in the direction $\mathbf{k}$ of laser propagation. The subsequent re-emission of photons is associated with momentum kicks; however, these are randomly distributed over the whole solid angle, so the net momentum gain over many events is zero.

Solving the optical Bloch equations in steady state for the two-level system gives [22]

$$\rho_{22} = \frac{\Omega^2/4}{\delta^2 + \Omega^2/2 + \Gamma^2/4}. \tag{2.2}$$

Here, $\Omega = d_{\text{eg}} E_0/\hbar$ is the Rabi frequency characterising the atom-field coupling strength with d_{eg} the electric dipole matrix element between ground and excited state, and $\delta = \omega - \omega_0 - \mathbf{k} \cdot \mathbf{v}$ is the detuning including the Doppler shift due to the atomic velocity $\mathbf{v}$. The force is maximised on resonance ($\delta = 0$) and saturates at $\hbar \mathbf{k} \, \Gamma/2$ for large intensities.

Molasses Cooling

The scattering force can be used to decelerate atoms using a counter-propagating beam, as shown in Figure 3(a). Because the momentum transfer of each absorbed photon is very small, this requires repeated cycling on a closed two-level transition, or repumper lasers to return atoms to the cooling cycle. Furthermore, only atoms within a certain velocity class can be addressed effectively, as they have to be resonant with the Doppler-shifted frequency of the light. For this slowing scheme, the light must be red-detuned ($\omega < \omega_0$), since the relative motion between the atom and the counter-propagating beam shifts the effective frequency up.

A second red-detuned laser beam from the opposite direction can be added to create a molasses configuration to cool the atoms to lower temperatures. In this case, atoms moving towards either laser are shifted closer to being on resonance with it, leading to increased scattering and a net damping force opposite to their motion, hence the label “molasses” cooling. The resulting velocity-dependent damping force is illustrated in Figure 3(b). This scheme can be extended to three dimensions by applying laser beams along all spatial directions. The final temperature of the atomic sample is bounded by two fundamental limits. The Doppler limit T_{D} arises from the balance between velocity selective cooling and random-walk heating due to spontaneous emis-

sion, and depends on the transition linewidth Γ. The recoil limit T_R is set by the irreducible momentum kick $\hbar k$ of a single scattered photon and the atom mass m. In the low-intensity limit, those are given in terms of temperature by

$$T_\mathrm{D} = \frac{\hbar\Gamma}{2k_\mathrm{B}} \qquad \text{and} \qquad T_\mathrm{R} = \frac{h^2}{mk_\mathrm{B}\lambda^2}. \tag{2.3}$$

Real atoms offer a range of possible transitions for optical molasses cooling, with an inherent tradeoff. Broader linewidths address a larger range of atomic velocities but are limited to a higher Doppler temperature, whereas narrower ones allow cooling to lower temperatures at the cost of capturing fewer atoms. Because of this, in practice both options are complementary and cooling schemes usually start with a broad transition to capture a high number of atoms, reducing the width of the velocity distribution before switching to a narrower linewidth to reach lower temperatures. The capture range of the narrow transition can be further extended using a comb of multiple frequency peaks, which bridges the two stages while keeping the temperature advantages of narrow-line cooling [23].

Magneto-Optical Traps

A magneto-optical trap (MOT) combines the effect of the scattering force with spatially varying transition frequencies [24]. This adds a position-dependent restoring force towards a common centre point on top of the velocity-dependent molasses force. A quadrupole magnetic field $\mathbf{B}(\mathbf{r}) = B'(x\,\hat{x} + y\,\hat{y} - 2z\,\hat{z})$ with B' the magnetic field gradient shifts the atomic sublevels. The Zeeman effect gives an energy splitting of magnetic sublevels m by

$$\Delta E = \mu_B g_J m B(\mathbf{r}) \tag{2.4}$$

where μ_B is the Bohr magneton and g_J the Landé g-factor [22]. Because B vanishes at the trap centre and increases with distance, the atomic transition frequency becomes position dependent. Counter-propagating red detuned beams with opposite circular polarizations drive $\Delta m = \pm 1$ transitions and the Zeeman shift changes the resonance condition differently for the two. As a result, an atom displaced from the trap centre is shifted closer to resonance with the beam that pushes it back towards the centre and further from resonance with the beam pushing it away. This imbalance produces the net restoring force towards the centre. Six beams are needed to have 3D confinement and cooling, which is again bounded by the Doppler limit.

A second method that uses the magnetic field induced shift of atomic transition frequencies is the Zeeman slower in which a spatially varying magnetic field is applied along the atomic beam. The resulting magnetic field compensates for the changing Doppler shift as the atoms are slowed, keeping the slowing laser resonant with the transition and enabling continuous deceleration.

2.1.2 The Dipole Force

While the MOT provides cooling and trapping, it relies on repeated photon scattering and fixes the atoms at the centre of the magnetic field. For coherent manipulation and flexible position control, such as the continuous guiding of atoms between two experimental positions, the optical dipole force provides a suitable conservative potential.

Arising from the interaction of the oscillating laser field with the laser-induced atomic dipole moment, it is a second-order perturbation effect. The underlying potential can be derived from the polarizability of the atom in the oscillating field. An electric field induces a dipole moment in the atom

$$\mathbf{d} = \alpha\, \mathbf{E}, \tag{2.5}$$

where α is the complex polarizability. Treating the atom as a classical oscillator driven at ω, it is given by [25]

$$\alpha = 6\pi\epsilon_0 c^3\, \frac{\Gamma/\omega_0^2}{\omega_0^2 - \omega^2 - i\,(\omega^3/\omega_0^2)\,\Gamma}. \tag{2.6}$$

The two parts of α have distinct physical roles. The imaginary part describes the absorptive, out-of-phase response that governs the photon scattering discussed in the previous section while the real part describes the dispersive, in-phase response of the atom and gives rise to the conservative dipole potential derived in the following. The imaginary part of α becomes negligible far away from resonance since the damping term $(\omega^3/\omega_0^2)\,\Gamma$ is small compared to $\omega_0^2 - \omega^2$. In this limit the interaction energy of the dipole with the field reduces to

$$U_{\mathrm{dip}}(\mathbf{r}) = -\frac{1}{2}\langle \mathbf{d}\cdot\mathbf{E}\rangle = -\frac{1}{2\epsilon_0 c}\,\mathrm{Re}(\alpha)\, I(\mathbf{r}) \tag{2.7}$$

where the cycle-averaged intensity is $I = \frac{1}{2}\epsilon_0 c\, E_0^2$ and an additional factor of $1/2$ arises from time-averaging the product $\mathbf{d}\cdot\mathbf{E}$ over one oscillation period [25]. For positive polarizability, the potential is attractive at intensity maxima, pulling atoms towards the focus of a Gaussian beam, which is the operating principle of an optical dipole trap. Atoms are trapped in the resulting dipole potential well when its depth exceeds their thermal kinetic energy $k_{\mathrm{B}}T$. For negative polarizability, atoms are repelled from intensity maxima instead.

The spatially varying potential can also be understood as a shift of the atomic energy levels induced by the laser field, known as the AC Stark shift or light shift. Since the energy of each atomic level is shifted by $U_{\mathrm{dip}}(\mathbf{r})$, any spatial variation in laser intensity translates directly into a position-dependent shift of atomic transition frequencies. This has important consequences in

experiments where atoms at different positions in a trap or lattice experience different resonance frequencies, which must be taken into account.

For a specific atomic level, the dipole force follows as the gradient of this potential,

$$\mathbf{F}_{\text{dip}}(\mathbf{r}) = -\nabla U_{\text{dip}}(\mathbf{r}) = \frac{1}{2\epsilon_0 c}\,\text{Re}(\alpha)\,\nabla I(\mathbf{r}). \tag{2.8}$$

Unlike the scattering force, the dipole force is conservative and does not operate via photon absorption-emission cycles. It is not fundamentally limited in magnitude and can be so large that it holds atoms against gravity. It is interesting to note different scaling with detuning, $U_{\text{dip}} \propto I/\delta$, while the scattering rate scales as I/δ^2. This means increasing the detuning at fixed trap depth suppresses unwanted photon scattering.

2.2 Gaussian Beams and Beam Shaping

From the previous section it is clear that shaping the intensity profile of the laser matters, as it determines the trapping potential and cooling efficiency. The propagation of most laser beams in our experiment can be modelled using Gaussian beams [26], where the beam radius w as a function of the propagation distance z is given by

$$w(z) = w_0\sqrt{1 + \left(\frac{z}{z_R}\right)^2}. \tag{2.9}$$

The Rayleigh range $z_R = {}^{\pi w_0^2}/_{\lambda}$ characterizes the divergence of the beam, and the waist w_0 is the minimum beam radius at the focus.

The radial intensity profile of a Gaussian beam at position z along the propagation axis, where r is the radial distance from the beam centre, is given by

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

where $I_0 = 2P/(\pi w_0^2)$ is the peak intensity and P is the laser power. The beam profiles along the propagation and radial directions are shown in Figure 4.

2.3 Intensity Noise and Heating

Fluctuations in the laser intensity affect cold atom experiments in several ways and decrease the quality of control over the atoms. For optical dipole traps, intensity noise translates directly into fluctuations of the trapping potential $U_{\text{dip}} \propto I$ which can heat trapped atoms and cause loss [27]. The dominant mechanism is parametric heating where intensity noise at twice the trap oscillation

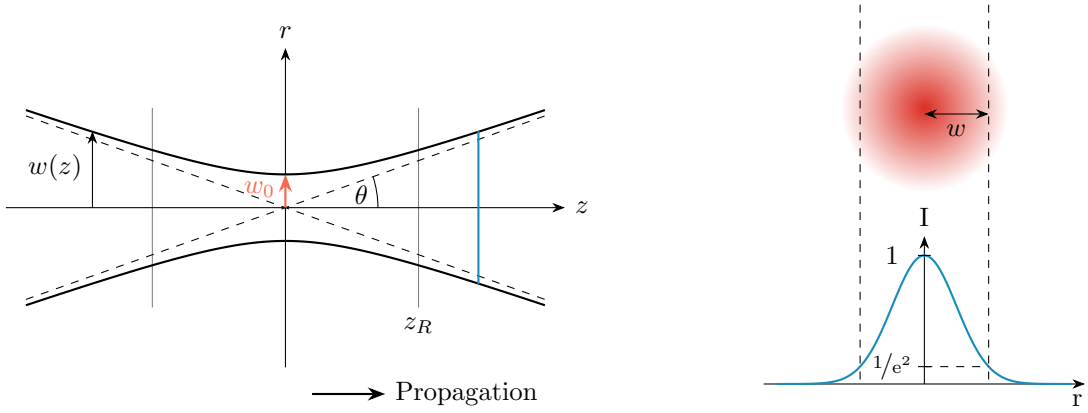


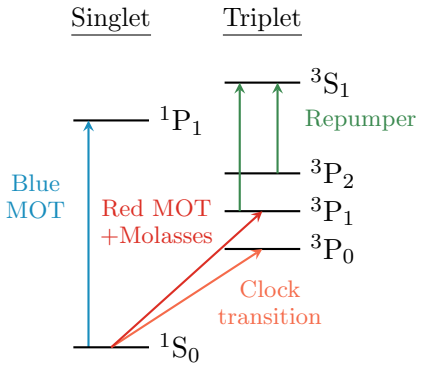
Figure 4: Gaussian beam profile showing beam parameters in the axial and transverse directions. The beam waist evolves axially according to Equation (2.9), the intensity profile in the transverse direction follows a Gaussian distribution given by Equation (2.10).

frequency (the frequency at which atoms oscillate in the trap potential) resonantly drives the motional modes and leads to exponential energy increase. Less critically, lower frequency noise can additionally cause random displacement of the trap centre. For molasses cooling, intensity fluctuations modify the scattering rate, however this is less crucial and usually not the limiting factor.

Beyond heating, intensity fluctuations lead to variations in the number of atoms loaded and transported and may hinder superradiant emission in this experiment. This will be discussed further in Chapter 5.

The relevant quantity is the power spectral density of the relative intensity fluctuations $\Delta I/I$, the relative intensity noise (RIN). Technical noise sources such as laser relaxation oscillations, acoustic perturbations, or electronic noise contribute to the RIN across a broad frequency range. Active intensity stabilization is therefore essential to ensure efficient and stable operation. The design and characterization of the stabilization system used in this experiment is discussed in Chapter 5.

3 | EXPERIMENTAL SETUP FOR CONTINUOUS SUPERRADIANCE ON ULTRACOLD ATOMS



Transition	f	$\Gamma/2\pi$	Ref.
$^1S_0 \to {}^1P_1$	460.7 nm	30.5 MHz	[28, 29]
$^1S_0 \to {}^3P_1$	689.3 nm	7.4 kHz	[28, 30]
$^1S_0 \to {}^3P_0$	698.4 nm	$\sim B$, $\lesssim 1$ mHz	[31]

Figure 5: Level diagram of strontium-88 with the relevant transitions for the experiment shown in Figure 6. The table shows the corresponding frequencies f and linewidths $\Gamma/2\pi$.

The experimental setup for achieving continuous superradiance on the mHz-linewidth clock transition in ultracold strontium follows the proposal of [32] and is schematically shown in Figure 6. It combines the continuous preparation of ultracold atoms in a MOT chamber with a transport scheme that continually injects µK atoms into a bowtie cavity in a science chamber designed for the realization of superradiance. The two chambers are spatially separated to limit the effects of stray light and magnetic fields from the MOT chamber onto the atoms during superradiant emission.

The relevant atomic transitions are shown in Figure 5. The dipole-allowed $^1S_0 \to {}^1P_1$ transition provides a broad linewidth suitable for initial cooling and capture. The spin-forbidden intercombination transition $^1S_0 \to {}^3P_1$ is much narrower and enables further cooling. The $^1S_0 \to {}^3P_0$ clock transition, additionally forbidden by angular momentum selection rules, is most interesting for achieving superradiance from a metrology perspective.

For both cooling transitions, the achievable temperature is bounded by the Doppler and recoil limits

$$T_\mathrm{D} = \frac{2\pi\hbar \cdot 30.5\,\mathrm{MHz}}{2k_\mathrm{B}} = \boxed{732\,\mu\mathrm{K}} \qquad T_\mathrm{R} = \frac{h^2}{88u\,k_\mathrm{B}\,(460.7\,\mathrm{nm})^2} = 1.03\,\mu\mathrm{K} \tag{3.1}$$

$$T_\mathrm{D} = \frac{2\pi\hbar \cdot 7.4\,\mathrm{kHz}}{2k_\mathrm{B}} = 0.18\,\mu\mathrm{K} \qquad T_\mathrm{R} = \frac{h^2}{88uk_\mathrm{B}\,(689.3\,\mathrm{nm})^2} = \boxed{0.46\,\mu\mathrm{K}}. \tag{3.2}$$

For the broad blue transition, the Doppler limit dominates, setting the achievable temperature. For the narrow red transition, the Doppler limit is below the recoil limit, so the recoil temperature becomes the fundamental limit of red MOT cooling.

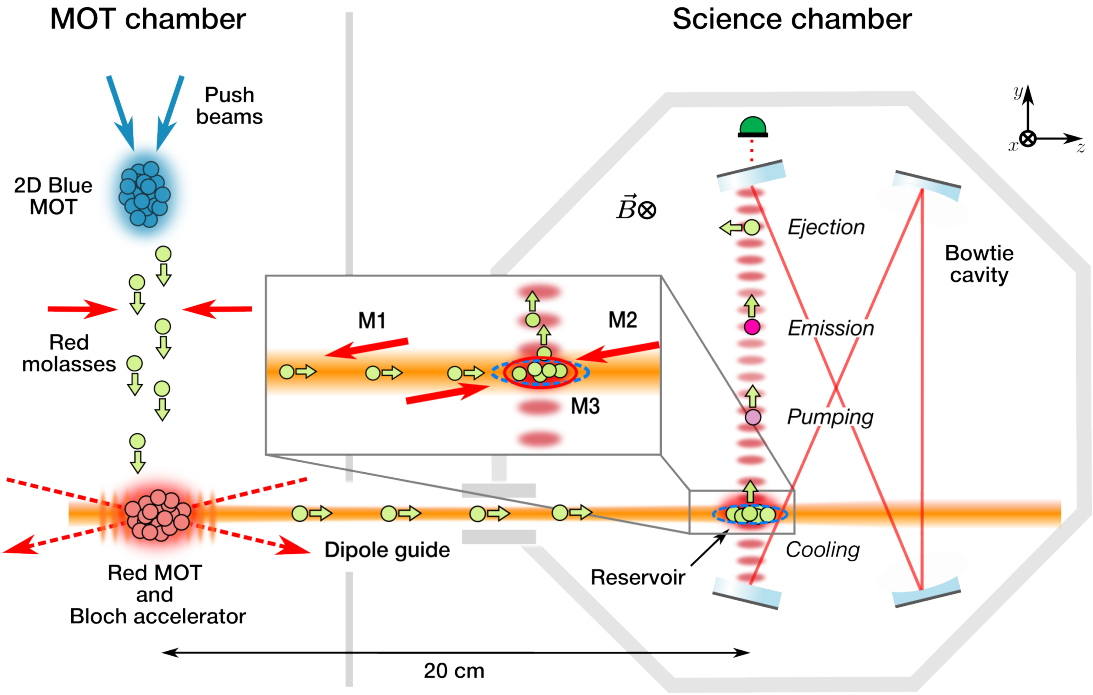


Figure 6: Schematic overview of the experimental apparatus for continuous superradiance with ultracold atoms. Atoms fall from a blue MOT into a red MOT where they are continuously coupled into a high-power optical dipole guide that transports them into the cavity. There, they are captured and cooled in a reservoir, before moving through the cavity in a moving optical lattice, where they are incoherently pumped into the excited state and emit via superradiance. Adapted from [32].

The experiment can be divided into different experimental steps the atoms undergo continuously but spatially separated. For the preparation of ultracold strontium, the atoms exit the oven and are transversely cooled before they pass a Zeeman slower and are captured in a blue MOT. This MOT consists of four horizontal MOT beams and two downward-pointing push beams that prevent atoms from escaping upward. This gives a cooling in the horizontal plane but lets the atoms leak out to the bottom, falling under gravity into a red 3D MOT below. Between the blue and the red MOT, red 2D molasses beams provide transverse cooling of the falling atomic cloud, collimating it before it enters the red MOT. The red MOT is a dual-position MOT combining a broadband (BB) and a single frequency (SF) component at different detunings following [33]. The BB stage captures the falling atoms, which are periodically transferred into the SF stage, where temperatures approaching the recoil limit of the red transition can be reached. The different detunings and powers of the BB and SF components determine their respective trap centers, resulting in spatially separated atomic clouds.

The µK temperature atoms can be loaded into a dipole guide that confines them radially, similar to a tube, and enables transport between the two chambers. The guide is formed by

a high-power, far-red-detuned infrared beam at 1070 nm wavelength, detuned sufficiently from all ground-state transitions to suppress inelastic scattering and associated losses. This beam is overlapped with the SF component of the red MOT, has a large waist and Rayleigh range, and is focused between the MOT and the cavity. This way, atoms slowly flow towards the focus under the conservative dipole force and decelerate again on the science-chamber side of the focus. However, because the axial intensity gradient is very small, this force is not sufficient to overcome the restoring force of the SF MOT. To nevertheless achieve loading of atoms into the dipole guide, a Bloch accelerator is overlapped with both the MOT and the dipole guide beam. It consists of two shallow angle beams with a slight frequency detuning between them, creating a moving optical lattice that accelerates the atoms in the dipole guide towards the science chamber [34].

Continuous superradiant emission requires a steady, uninterrupted supply of atoms through the cavity. Achieving this requires a mechanism that can transport atoms into and through the cavity while keeping them cold and trapped. A moving optical lattice along the cavity axis provides exactly this. To realize it, the atoms are first captured in a reservoir trap formed by a second, crossed dipole trap, where they are cooled using the molasses beams M1-M3 shown in Figure 6. This stops their transverse motion and prepares them for slow transport along the perpendicular cavity axis. The bowtie ring cavity allows for creating such a moving optical lattice from two counter-propagating beams with a small relative frequency offset. For sufficiently cold atoms and slow lattice velocities the resulting dipole potential traps atoms in the lattice antinodes.

They are then pumped into the excited state, can synchronize their dipoles through the cavity-mediated interaction and emit collectively through superradiance. It is essential that the pumping is incoherent, which means the laser does not imprint any phase on the atoms, which could otherwise transfer its frequency instability to the emitted light. One approach is to pump atoms into the higher lying $^3S_1$ state, from which they spontaneously decay into the $^3P_0$ clock state with a random phase, a scheme presented in [32]. Because atoms can also decay from $^3S_1$ into the $^3P_1$ and $^3P_2$ states, repumpers are needed to bring these atoms back into the clock state.

The apparatus can host different strontium isotopes, but the choice of isotope determines both the scheme for operating the clock transition and the achievable number of atoms. The fermionic isotope ^{87}Sr has a nonzero nuclear spin, whose hyperfine interaction weakly allows the otherwise forbidden clock transition, giving it a natural linewidth on the mHz level and making it directly accessible. The bosonic isotope ^{88}Sr has no hyperfine structure and the $^1S_0 \rightarrow {}^3P_0$ clock transition is doubly forbidden, with no coupling to light at zero field. It can only be opened by applying a strong magnetic field of several hundred gauss to mix in a small fraction of the $^3P_1$ state [32, 35]. Without this field there is no clock transition to drive, and hence no mHz superradiance in ^{88}Sr. This strong, controlled field must be maintained in the science chamber,

where it is incompatible with the MOT field gradients, a crucial reason why the cooling and science regions are kept in separate chambers. Implementing it with the required precision and magnetic isolation between the two regions is experimentally challenging. In terms of atom number, the bosonic isotope has a higher natural abundance (82.6 %) than the fermionic isotope (7.0 %), which increases the achievable superradiant output for ^{88}Sr [36].

4 | TOWARDS CONTINUOUS TRANSPORT OF ULTRACOLD STRONTIUM

This chapter presents the different stages involved in preparing and continuously transporting ultracold atoms between the MOT chamber and the science chamber, and describes several experimental steps implemented to achieve this.

4.1 Blue and Red MOT Cooling to µK Temperatures

To have efficient control over the atoms using optical potentials, it is crucial that they are prepared in low kinetic energy states. This section describes how the magneto optical traps are optimized to reach this regime and investigates their dynamics in the context of the planned continuous transport scheme.

As described in Chapter 3, the atoms fall from the 2D blue MOT into a red MOT operating on the narrow intercombination line, combining broadband (BB) and single frequency (SF) light. This follows the general idea outlined in Section 2.1.1, where a broad linewidth stage prepares atoms for narrow-line cooling. What is special here is the spatial rather than temporal separation of the two stages. Operated at different detunings, the BB and SF components form spatially separated clouds. The BB stage runs continuously and captures the falling atoms, while the SF stage is pulsed to periodically transfer atoms. Because this switching is fast compared to the loading and cooling dynamics, the SF MOT prepares a quasi-continuous source of sub-µK atoms [33].

The blue MOT is monitored via fluorescence imaging, as it involves fewer parameters and is comparatively straightforward to optimize. The red MOT can be imaged in two ways, both along the horizontal x-direction in Figure 6. Continuous red fluorescence imaging can be used for rough alignment, while absorption imaging on the blue transition is used to extract the atom number and to determine the temperature, position and widths of the cloud. Especially in dense MOTs, as in our experiment, absorption imaging is more reliable for extracting atom number, as fluorescence imaging can underestimate atom number due to reabsorption of emitted photons by neighboring atoms. In addition, pulsed loading sequences are useful here, as imaging before the MOT reaches steady state avoids the high optical depths that would otherwise saturate the absorption signal.

The large velocity capture range of the BB stage introduced above is achieved by spectrally broadening the red cooling light into a frequency comb. This comb, used for both the red MOT and the 2D molasses beams, is generated by periodically ramping the output frequency of a direct digital synthesizer (DDS) that drives the acousto-optic modulator (AOM) via a radiofrequency (RF) signal to control the laser frequency. Because the DDS frequency varies periodically with

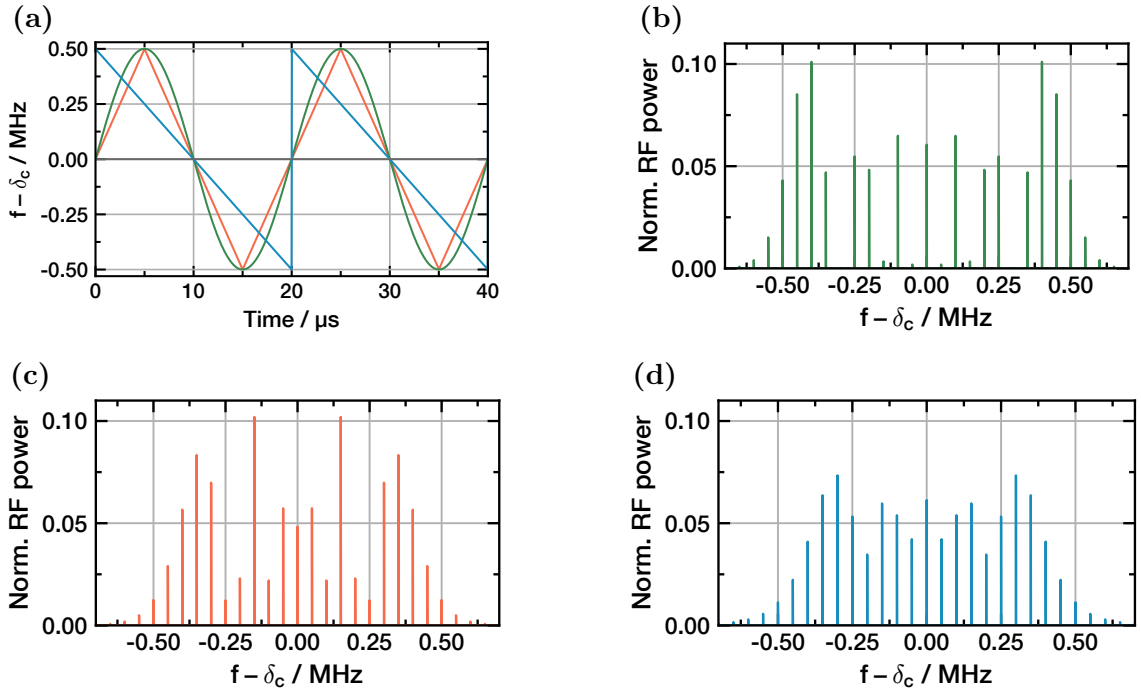


Figure 7: Broadband cooling method. (a) Frequency modulation of the AOM drive signal over two periods for sine (green), triangle (orange) and sawtooth (blue) function, centred around the detuning δ_c with half-width $\Delta = 0.5\,\mathrm{MHz}$ and modulation frequency $f_\mathrm{m} = 50\,\mathrm{kHz}$. (b)-(d) Resulting RF power spectrum for sine, triangle and sawtooth modulation respectively, showing frequency combs of lines spaced by f_m within the modulation envelope of half-width Δ. Spectra are normalized to their total power.

modulation frequency f_m, its power spectrum develops a discrete comb structure with lines spaced by f_m. The different comb lines address atoms with different velocities, extending the effective capture range of the MOT. The parameters of this ramp are the central detuning δ_c, the half-width of frequencies that are scanned Δ, and the modulation frequency f_m with which the ramp in time is implemented.

The shape of the frequency ramp determines the spectral envelope of the comb. Triangular, sinusoidal and sawtooth-shaped ramps and the resulting spectra are shown in Figure 7. The sinusoidal modulation produces high power components at large detuning and has many central frequency components of very low power. This is because the sinusoidal ramp spends more time near its turning points at $\pm\Delta$ than at low detuning, whereas the triangular ramp sweeps at constant rate, resulting in stronger comb teeth near the line centre. The sawtooth modulation gives a spectrum with power more evenly distributed among the frequency components. In this experiment a triangular ramp was chosen. However, it may be useful to test the different

modulation waveforms on atoms and compare their cooling performance at the same central detuning, half-width and modulation frequency.

Optimization of the Cooling System

To optimize the red MOT, the frequencies of the relevant laser beams were scanned and set to the value maximizing the atom number measured in absorption imaging. First, the frequencies of the blue imaging beam and the blue MOT beam were optimized. Although the latter is a parameter of the blue MOT, it was optimized using the red MOT atom number, as maximizing the atom number arriving at the red MOT is the relevant goal and absorption imaging at this stage provides a more reliable signal. Because the later red narrow-line stages capture colder atoms more efficiently, optimizing on the red MOT atom number favors a detuning near the Doppler-cooling optimum of the blue MOT $\delta = -\Gamma/2 = -2\pi \cdot 15.3\,\mathrm{MHz}$ [22], rather than a larger detuning that would maximize blue MOT capture. The measured value of $\delta/2\pi = -14.1\,\mathrm{MHz}$ agrees well with this.

For the red beams, the start, stop, and modulation frequencies were optimized for the 2D molasses, the horizontal MOT beams and the vertical MOT beam. The vertical MOT beam is different because it holds the atoms against gravity, and is not retroreflected at the top. The resulting parameters are listed in Table 1. Even though only the atom number was optimized, this simultaneously reduced the temperature of the cloud to a point that is sufficient for the BB MOT. The alignment of the beams and the waveplate setting the polarization of the vertical beam were adjusted to maximize the atom number, while the waveplates of the horizontal beams were optimized for low temperature, because in this case optimizing for atom number increased the temperature. The temperature was determined from time-of-flight measurements discussed below.

The optimized MOT is shown in Figure 8(a), and the single frequency MOT obtained after the ramp sequence of [33] is shown in Figure 8(b). This ramp sequence is designed to transfer atoms from the BB to the SF MOT with little losses. First, the BB intensity is ramped down while simultaneously ramping up the SF intensity, ensuring the atoms remain continuously

Table 1: Parameters optimized in frequency scan to maximize red MOT atom number. Central detuning δ_c, half width frequency span Δ and modulation frequency f_m for red broadband MOT as shown in Figure 7.

	δ_c / MHz	Δ / MHz	f_m / MHz
2D Molasses	−0.42	±0.41	0.07
Red MOT horiz.	−1.72	±0.68	0.06
Red MOT vert.	−1.54	±0.47	0.03
Blue MOT	−14.1	—	—

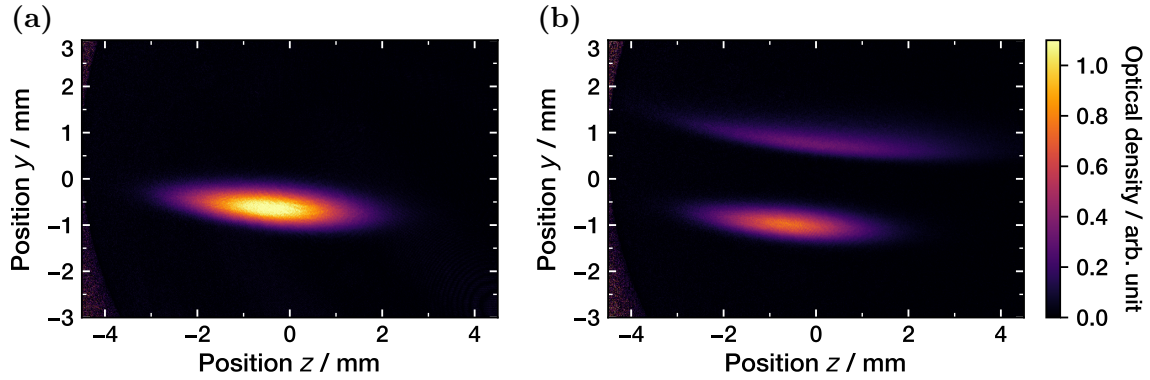


Figure 8: Absorption images of the red MOT optical density. (a) Broadband MOT and (b) Broadband + single-frequency MOT after applying an intensity ramp sequence as described in [33].

confined. A subsequent ramping down of both components allows the cloud to cool further and reduces the photon scattering rate. The magnetic field gradient, set by the anti-Helmholtz coil current, was optimized to around $0.6\,\mathrm{G\,cm^{-1}}$. To account for changes in the background magnetic field and to control the position of the MOT, magnetic bias fields in each direction are generated using Helmholtz coils. Because the MOT is so sensitive to magnetic field changes, the coil currents must be reoptimized regularly, especially when changes in the experimental environment affect the local magnetic field.

Red MOT Dynamics

For the characterization of the red MOT, Figure 9 shows measurements and analysis of the properties relevant for continuous transport. To understand the limits of continuous loading and transport, it is interesting to look at the loading curve and the lifetime of the MOT. Both are described by the general rate equation for N_{MOT} the number of atoms in the MOT

$$\dot{N}_{\mathrm{MOT}} = L - N_{\mathrm{MOT}}/\tau, \tag{4.1}$$

where L is the loading rate and τ is the mean survival time of an atom in the trap. Integrating Equation (4.1) with the initial condition $N(0) = 0$ gives the loading curve

$$N_{\mathrm{MOT}}(t) = N_{\mathrm{max}} \cdot \left(1 - \mathrm{e}^{-t/\tau_{\mathrm{L}}}\right), \tag{4.2}$$

where $N_{\mathrm{max}} = L\tau$. Fitting this expression to the data in Figure 9(a) yields the loading time constant τ_{L} and a saturation atom number of 3.7×10^8 atoms in the red BB MOT. The loading at early times is approximately linear, which motivates fitting a linear equation $N_{\mathrm{MOT}}(t) = L_0 t$

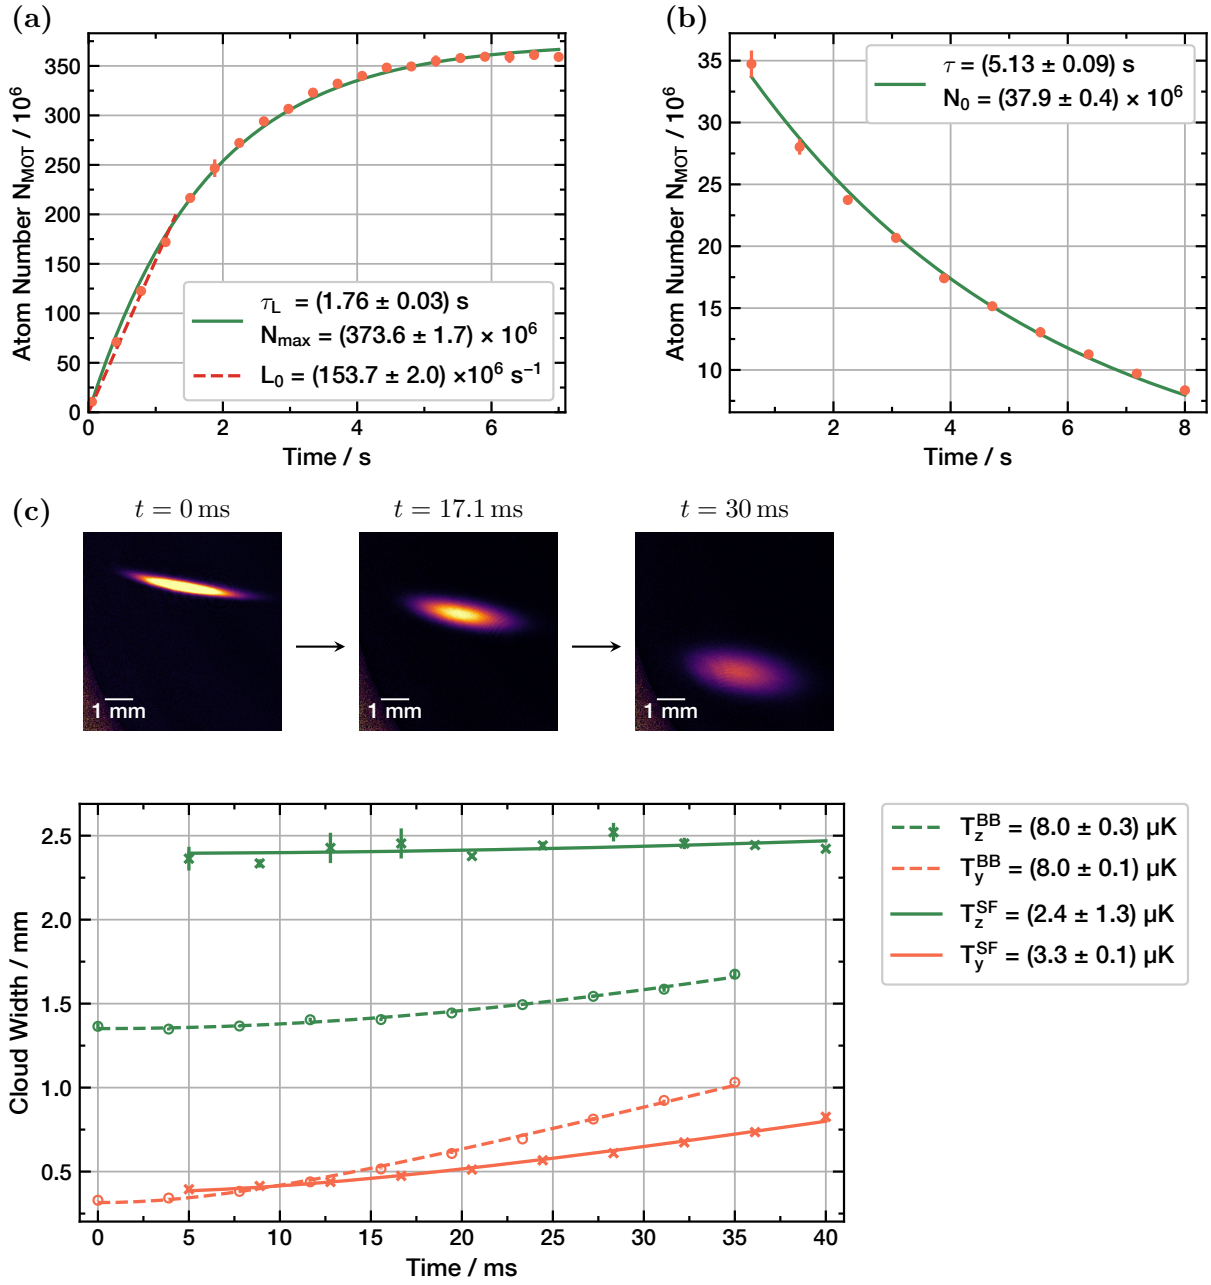


Figure 9: Red MOT characterization. Each data point represents the mean of three repeated measurements, with error bars indicating the standard deviation. (a) BB MOT loading curve fitted with Equation (4.2), with a linear fit to the initial loading rate for $t < 1.15\,\mathrm{s}$. (b) BB MOT lifetime curve fitted with Equation (4.3). (c) Time-of-flight (TOF) measurement of the BB and SF MOT cloud widths, with absorption images of the BB MOT illustrating the cloud expansion and free fall. Dashed lines indicate the BB MOT fit, solid lines the SF MOT fit, both using Equation (4.4).

that directly measures the initial loading rate L_0. For the first 1.15 s this yields $L_0 = 153.7(20) \times 10^6\,\mathrm{s}^{-1}$. This value gives a representative atom flux available for loading under our conditions, depending on the oven temperature and on the current technical conditions such as the laser frequency locks and beam alignment. For the lifetime measurement, the loading rate is set to zero and the initial atom number $N(0) = N_0$, so that Equation (4.1) reduces to the simple exponential decay

$$N_{\mathrm{MOT}}(t) = N_0 \cdot \mathrm{e}^{-t/\tau}. \tag{4.3}$$

The fit shown in Figure 9(b) gives a lifetime of $\tau = 5.13(9)\,\mathrm{s}$.

The temperature of the atomic cloud is extracted using the standard time-of-flight (TOF) method which relies on the fact that the cloud's temperature is directly linked to the kinetic energy, and therefore the velocity distribution of the atoms. After release from the trap, the cloud expands ballistically during free fall with higher temperature meaning faster expansion which can be seen in the insets in Figure 9(c). The temperature T can be extracted from the cloud width $\sigma(t)$ via [37]

$$\sigma(t) = \sqrt{\sigma_0^2 + k_{\mathrm{B}} T t^2 / m} \tag{4.4}$$

where σ_0 is the initial width of the MOT and m is the atom mass. The cloud width is measured separately in the vertical and horizontal directions, as shown in Figure 9(c). The fit yields 8.0 µK for the broadband MOT. The horizontal width of the SF MOT shows considerably larger error bars, since the cloud becomes strongly elongated in this configuration and the extracted size is then more sensitive to fluctuations in the cloud shape. The temperature of the SF MOT should therefore be extracted from the vertical expansion only and in this case is 3.3 µK. The SF MOT parameters have not been fully optimized for this measurement, and temperatures below 1 µK are achievable [33].

4.2 Optical Dipole Guide and Atomic Transport

To transport the ultracold atoms from the red MOT position to the cavity, they are loaded into a weakly focused, far-detuned laser beam. In contrast to a tightly focused optical tweezer or a dipole trap optimized for tight confinement, the dipole guide is designed to maintain a large beam diameter over the entire transport distance. This reduces the pull towards the horizontal focus while keeping the radial trapping force strong enough to support the atoms against gravity throughout. This section theoretically develops the requirements of this beam by giving a calculation of the effective potential depth. It further explains the experimental particularities of setting up the optics and should give some practical insights on working with high power lasers.

Theoretical Considerations and Power Requirements

The dipole potential experienced by a ground-state ^{88}Sr atom in the guide is given by Equation (2.7), with the intensity profile of Equation (2.10) and a beam radius that evolves along the transport axis z according to Equation (2.9). We use a beam waist of $w_0 = 140\,\mu\text{m}$, focused at the midpoint between the MOT and the cavity. The correspondingly large Rayleigh range $z_\text{R} = 5.76\,\text{cm}$ provides strong transverse confinement while forming only a shallow axial minimum at the center of the transport region, so that atoms displaced along the axis are gently pulled toward it. At the guide wavelength $\lambda = 1070\,\text{nm}$, the ground-state polarizability of ^{88}Sr is $\alpha_\text{au} = 240\,\text{a.u.}$ [38], converted to SI units via $\alpha_\text{SI} = 4\pi\epsilon_0 a_0^3\,\alpha_\text{au}$, where a_0 is the Bohr radius [39].

In the coordinate system of Figure 6, the guide propagates along z and gravity acts along y. The reservoir beam propagates along x and crosses the guide at the cavity position, 10 cm beyond the focus.

The effective trap depth differs from the bare optical depth because gravity pulls the atoms slightly below the beam axis ($y = 0$). The total potential at each point is

$$U_\text{tot}(x, y, z) = U_\text{dipole}(x, y, z) + U_\text{reservoir}(x, y, z) + mgy. \tag{4.5}$$

Figure 10(a) shows vertical cuts of the potential at the two positions of strongest (guide focus, $z = 0$) and weakest (MOT position, $z = -100\,\text{mm}$) confinement, comparing the purely optical potential with the total potential including gravity. Gravity tilts the well, shifts the equilibrium downward, and lowers the barrier on the downhill side. The shift grows as the beam widens away from the focus. The atoms settle at the vertical equilibrium where the total potential is stationary,

$$\left.\frac{\partial U_\text{tot}}{\partial y}\right|_{y=y_\text{eq}(z)} = 0, \tag{4.6}$$

evaluated at each axial position z, with $x = 0$. This condition has two solutions: a minimum, at which the atoms are trapped, and a maximum below the beam axis. The energy difference between those is the effective barrier ΔU that an atom must overcome to escape the guide, also shown in Figure 10(a). The equilibrium sits visibly below $y = 0$, increasingly so away from the focus, where the beam widens and the vertical confinement weakens.

Away from the cavity, where the reservoir contribution is negligible, Equation (4.6) reduces to a balance between the guide's optical restoring force and gravity, which can be written explicitly as

$$\frac{2\alpha I_0 w_0^2}{\epsilon_0 c\, w(z)^4}\, y_\text{eq}\, \exp\left(-\frac{2y_\text{eq}^2}{w(z)^2}\right) = mg. \tag{4.7}$$

This is transcendental in y_eq and is solved numerically at each z, in general the reservoir term is included and the stationary points of U_tot are determined. Figure 10(b) shows, from top

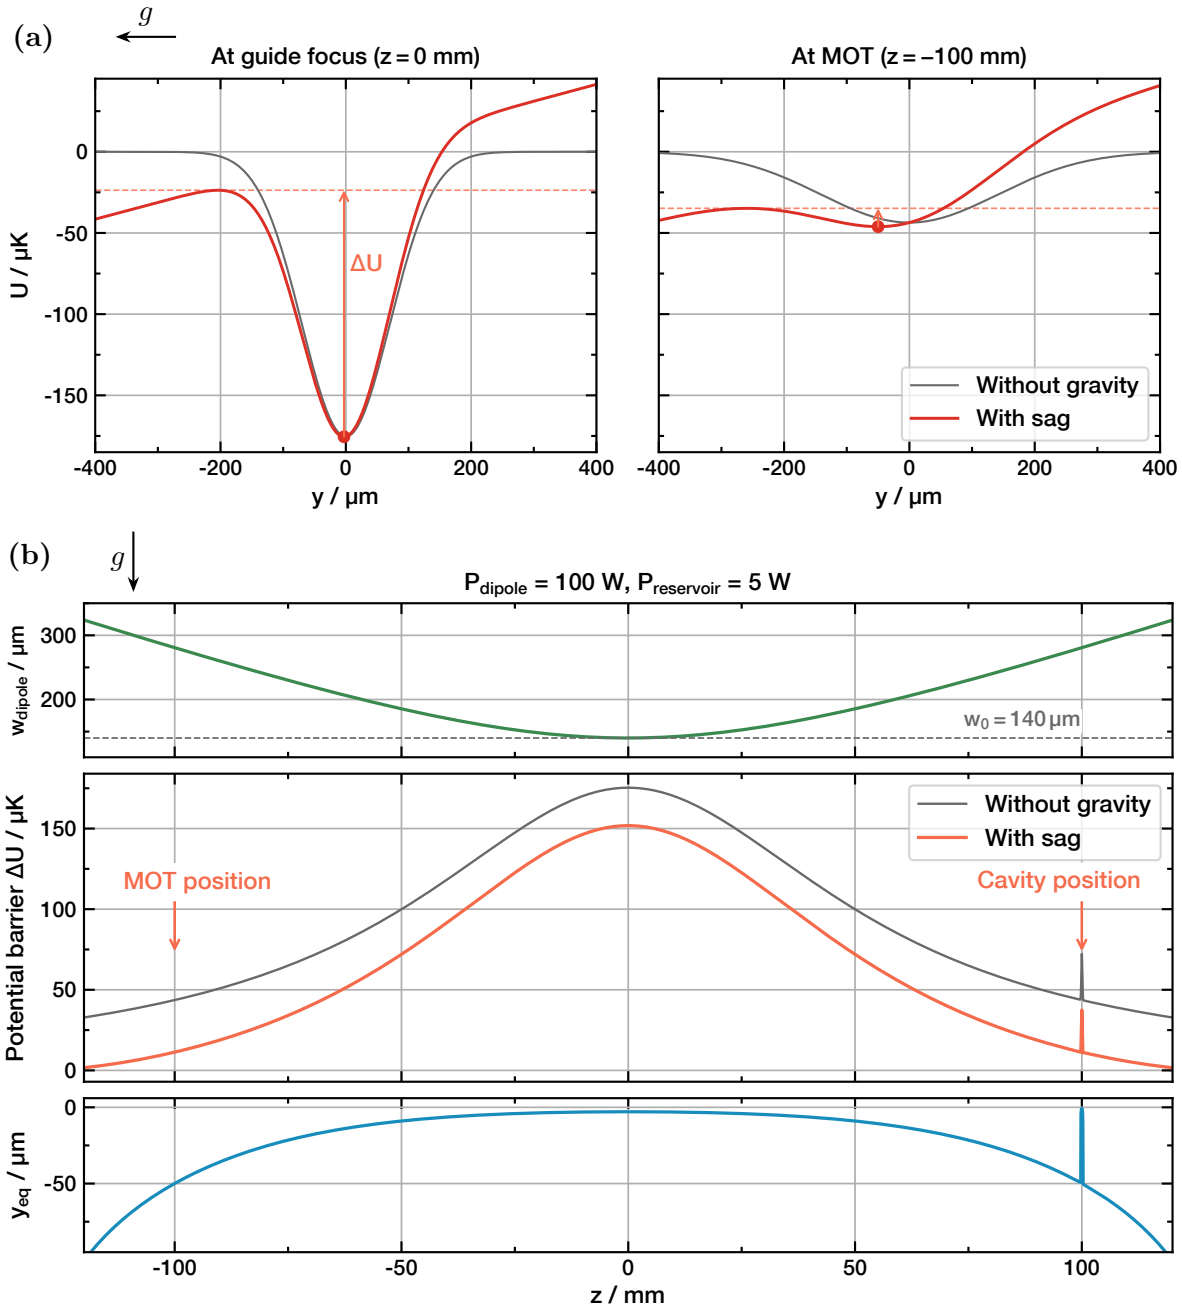


Figure 10: Theoretical calculation for the dipole guide, using the coordinate system defined in Figure 6. (a) shows the effect of the gravitational potential in the vertical direction, which lowers the effective potential barrier that holds the atoms in the guide. (b) shows the physics along the dipole guide's propagation axis, from the MOT to the cavity position: the beam radius of the guide with a weak focus of 140 μm (top), the trapping barrier with and without gravitational sag (middle), and the equilibrium position y_{eq} to which the atoms sag, measured from the beam center $y = 0$ (bottom). The sharp feature near $z = 100\,\mathrm{mm}$ marks the cavity position, where the crossed reservoir beam intersects the guide.

to bottom, the guide radius, the effective barrier with and without gravitational sag, and the equilibrium sag position y_{eq} along the transport axis, for a guide power of 100 W and a reservoir power of 5 W. The reservoir profile is described in the following section.

The effective barrier sets an upper bound on the kinetic energy an atom can have before escaping the guide. For reliable loading and to tolerate imperfections such as an inhomogeneous beam profile, it should exceed the thermal energy by at least a factor of two to three. The barrier is deepest at the focus and shallowest at the ends of the transport region. At a guide power of 100 W it drops to 11 µK at the furthest positions $|z| = 100\,\mathrm{mm}$, where the MOT and the reservoir trap additionally support the atoms. Since the transported cloud prepared in the single-frequency MOT has a temperature of only up to a few microkelvin, this barrier still exceeds the thermal energy by more than a factor of three, so the guide is expected to confine the atoms. There should be an additional margin of power to adjust this in the experiment if the transport or loading efficiency proves too low.

Because the barrier is limited by the largest beam radius along the path, at the MOT and cavity positions, the choice of waist w_0 directly affects the power requirement. For a fixed distance L from the focus to the far end of the transport region, the radius there, $w(L) = \sqrt{w_0^2 + (L\lambda/\pi w_0)^2}$, is minimized for

$$w_0^{\mathrm{opt}} = \sqrt{\frac{L\lambda}{\pi}} \;\Leftrightarrow\; z_R = L, \tag{4.8}$$

giving $w_0^{\mathrm{opt}} = 185\,\mu\mathrm{m}$ for $L = 100\,\mathrm{mm}$ and $\lambda = 1070\,\mathrm{nm}$. The waist $w_0 = 140\,\mu\mathrm{m}$ assumed above is therefore slightly tighter than this optimum.

Experimental Implementation

To meet the requirements of stable trapping and support against gravity over a long distance, a high-power IPG fiber laser with a maximum output of 400 W at $\lambda \approx 1070\,\mathrm{nm}$ is used. An industrial laser of this kind is well suited, since far off-resonant trapping requires neither a narrow linewidth nor high coherence. A partially incoherent source can in fact be preferable, as it suppresses interference fringes from reflections at optical surfaces and vacuum windows.

The light of the IPG fiber laser is delivered by a fiber onto a breadboard next to the MOT chamber, with the optical layout shown in Figure 11. A practical challenge is that the minimum stable output power of the laser is approximately 25 W, far more than is convenient for beam shaping and alignment. Most of this power is therefore split off with a waveplate and PBS and directed into a beam dump.

Working with only the minimum transmission through the PBS, about 300 mW, proved impractical. In this configuration the beam size out of the laser differed from what was expected

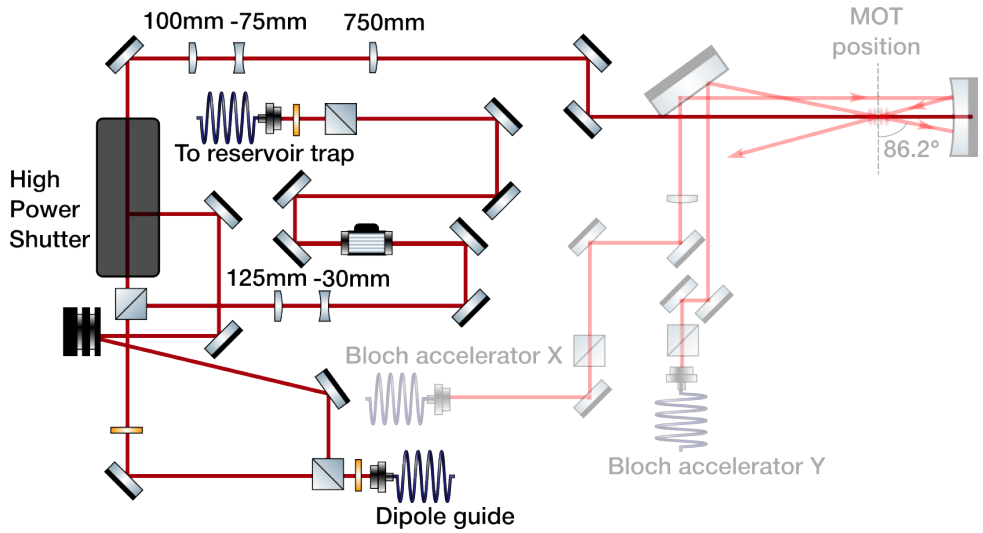


Figure 11: Optical setup of the dipole trap optics and Bloch accelerator beams. The latter are shown semi-transparent, as they are not discussed here. The light of the IPG fiber laser is delivered to the breadboard through a fiber. Its power is split at a waveplate and PBS, and one part is sent through an AOM and a fiber that goes to the science chamber breadboard in Figure 13 where it is used to create the reservoir beam. Symbols as defined in the List of Optical Elements (Figure 1).

and clear wavefront aberrations appeared. Focusing with a lens produced ring-shaped artifacts in the beam profile and a mode that deviated noticeably from Gaussian. The beam quality improved at higher power, so alignment was first performed using the leakage transmitted through a high-reflectivity mirror, where no aberrations were found. Here, no rings appeared after the lens and the AOM diffraction efficiency was high. The setup was then switched back to the waveplate and PBS combination and realigned to correct for the slight beam displacement introduced by the mirror transmission. That the mode degraded specifically at minimum PBS transmission points to the finite polarization purity of the laser as the cause. The transmitted light in this configuration is dominated by the small, orthogonally polarized component, which likely carries a distorted mode.

Since an AOM cannot be used at powers of potentially more than 100 W, the beam is switched with a high-power mechanical shutter that reflects it out of the path and into a water-cooled beam dump. All optical elements along the main path, including this shutter, must be compatible with the beam power requirements, including high-reflectivity mirrors, low-absorption UV fused silica lenses, and high-power zero-order waveplates.

After the shutter, the guide beam is shaped by a telescope with focal lengths of 100 mm and −75 mm that reduces the beam size. The telescope is followed by a 750 mm lens placed one focal length from the midpoint of the transport region. This lens sits in a cage mounted on

a z-translation stage, so that the focus position can be adjusted along the beam axis without significantly disturbing the alignment.

Should the beam prove too tightly focused, the telescope can be changed to give a looser focus with a Rayleigh range of 0.1 m. Replacing the 100/−75 mm telescope with a 100/−60 mm pair yields a waist of about 183 µm, close to the optimum derived above. The current beam shaping was planned with an earlier version of the calculation, which has since been refined.

On the same breadboard, before the shutter, a fraction of the beam is split off and sent through an AOM and a fiber to form the reservoir trap, delivering about 5 W, set by the fiber's power limit. The Bloch accelerator beams are launched from the same breadboard. The Bloch accelerator consists of two slightly detuned beams, one of them retro-reflected inside the chamber by a mirror with a hole through which the atoms and the dipole guide pass. The beams overlap at the MOT position, nearly counter-propagating and crossing at a shallow angle, and interfere to form a moving optical lattice. Its velocity is set by the small frequency difference between the two beams, so that the lattice accelerates the atoms towards the cavity.

Loading Atoms into the Dipole Guide

Initial loading of atoms into the dipole guide had already been demonstrated in the group. Here, this loading was further optimized. The guide focus is placed at the red MOT position by translating the final lens in its cage. This maximizes the local intensity, and hence the trap depth, exactly where the atoms are prepared in the MOT. Concentrating the available power at this point gives the deepest possible initial trap and reduces the power needed for loading, which was about 25 W here. Shifting the final lens by 10 cm can later move the focus to the midpoint between the MOT and the cavity for transport.

The magnetic bias fields are adjusted to overlap the MOT with the dipole guide and in the loading sequence the dipole guide is turned on before loading the MOT. Figure 12 shows an absorption image with atoms in the dipole guide. The MOT cloud has expanded for 25 µs and fallen under gravity, while the atoms held in the dipole guide remain trapped. The loading efficiency here is around 2 %, optimized by scanning the bias magnetic fields and the timing of the loading sequence. A colder, smaller SF cloud and more precise spatial alignment can increase the number of atoms within the dipole guide trapping volume. In addition, the dipole guide induces a spatially varying AC Stark shift on the intercombination line used for the MOT, so adjusting the SF MOT frequency to compensate this light shift at the guide position restores optimal cooling where the atoms are loaded.

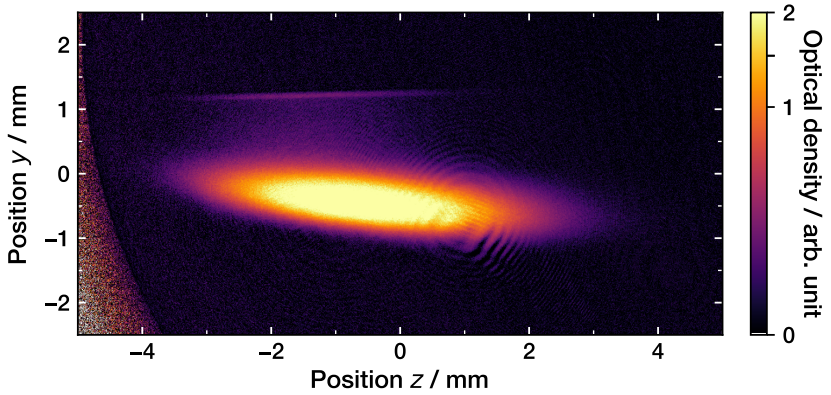


Figure 12: Absorption image after a dipole guide loading sequence with the expanded MOT and the atoms captured in the dipole guide. The color scale is nonlinear so that the weaker signal from the guided atoms is visible alongside the much brighter MOT cloud.

4.3 Slowing and Recapture in Reservoir Trap

Since the dipole guide is a conservative potential, atoms passing it arrive at the reservoir position with approximately the same velocity at which they entered the guide at the MOT, accelerated with the Bloch beams. Simulations in [32] predict a velocity of approximately $0.1\,\mathrm{m\,s^{-1}}$ at the reservoir position. To recapture atoms at this velocity, a sequence of broadband cooling beams is applied at the reservoir area, as illustrated in Figure 6.

The cooling sequence consists of three 689 nm beams which all use the broadband frequency modulation described in Section 4.1. Beam M1 is a unidirectional beam directed against the atom propagation direction, providing a decelerating scattering force over a broad velocity range. Beams M2 and M3 are retroreflected, forming a molasses configuration in both horizontal directions to cool transverse motion.

The central detuning δ_c of M1 is chosen to address atoms arriving at $v \approx 0.1\,\mathrm{m\,s^{-1}}$, where the Doppler resonance condition gives

$$\delta_c = -kv = -\frac{2\pi}{\lambda}\, v = -2\pi \cdot 145\,\mathrm{kHz}. \tag{4.9}$$

The required broadband half-width Δ is set by the velocity spread of the incoming atoms: assuming a Gaussian spread of $\pm 1\,\mathrm{cm\,s^{-1}}$ corresponds to a frequency spread of $\pm 1\,\mathrm{cm\,s^{-1}}/689\,\mathrm{nm} = \pm 14.5\,\mathrm{kHz}$ which is needed to capture the full distribution. Final parameters will be optimized directly on atoms.

The reservoir trap is a crossed optical dipole trap positioned at the cavity axis and spatially overlapped with the dipole guide beam. Additional beams directed toward the science chamber are set up on the same optical breadboards, including optical pumping and imaging beams.

It should be noted that the high-power dipole guide and reservoir beams introduce a spatially varying AC Stark shift on the cooling transition (Section 2.1.2). The central detunings of the molasses beams must therefore be adjusted to compensate this shift at the reservoir position, so that the cooling remains resonant.

Beam shaping

The beam paths for all beams directed at the science chamber are set up on three optical breadboards surrounding the chamber, shown in Figure 13. Blue absorption imaging perpendicular to both the cavity axis and the dipole guide is implemented using a large beam, with a camera and focusing system mounted below the front breadboard. The optics for coupling light into the cavity, for cavity locking and probing superradiance, are installed on a separate breadboard on top of the science chamber.

All molasses and the pumping beam require a waist of approximately $w_0 = 300\,\mu\text{m}$ at the cavity, matched to the spatial extent of the atom cloud at the reservoir position. Since the beams enter the science chamber from different distances and angles, each path requires individually designed focusing optics. The long distance between the last optical element and the cavity explains why telescopes and long focal-length lenses are chosen to achieve the required waist. Beams M1 and M2 are launched from the right side of the chamber, while M3 enters perpendicularly from behind. M1 and M2 enter the chamber with a slight downward component, because the last mirror before the science chamber sits above the dipole-guide beam, which exits the chamber after passing through the cavity. The pumping beam and the molasses beams M2 and M3 are retroreflected on the front breadboard, forming the molasses configuration. The retroreflection optics are designed to reproduce the beam waist at the reservoir position (Section 4.4.2). M3 and the pumping beam exit the chamber towards the front breadboard, while M2 is reflected on an angled mirror in the chamber with a central hole, which allows the atoms to pass through as they travel into the science chamber while reflecting the beam to the front breadboard. The molasses and pumping beams sharing the same spatial path are combined using d-shaped mirrors. Multiple beams of different wavelengths are combined and separated using dichroic mirrors, specified in Figure 13.

The beam profile at each position is characterised by fitting a Gaussian to the camera image of the beam according to Equation (2.10) which yields $w(z)$. Propagation through the optical system is modelled using the ABCD ray transfer matrix formalism [40]. Within the paraxial and thin-lens approximations, each optical element is represented by a 2×2 matrix, and the beam waist at any position is obtained by propagating the complex beam parameter q through

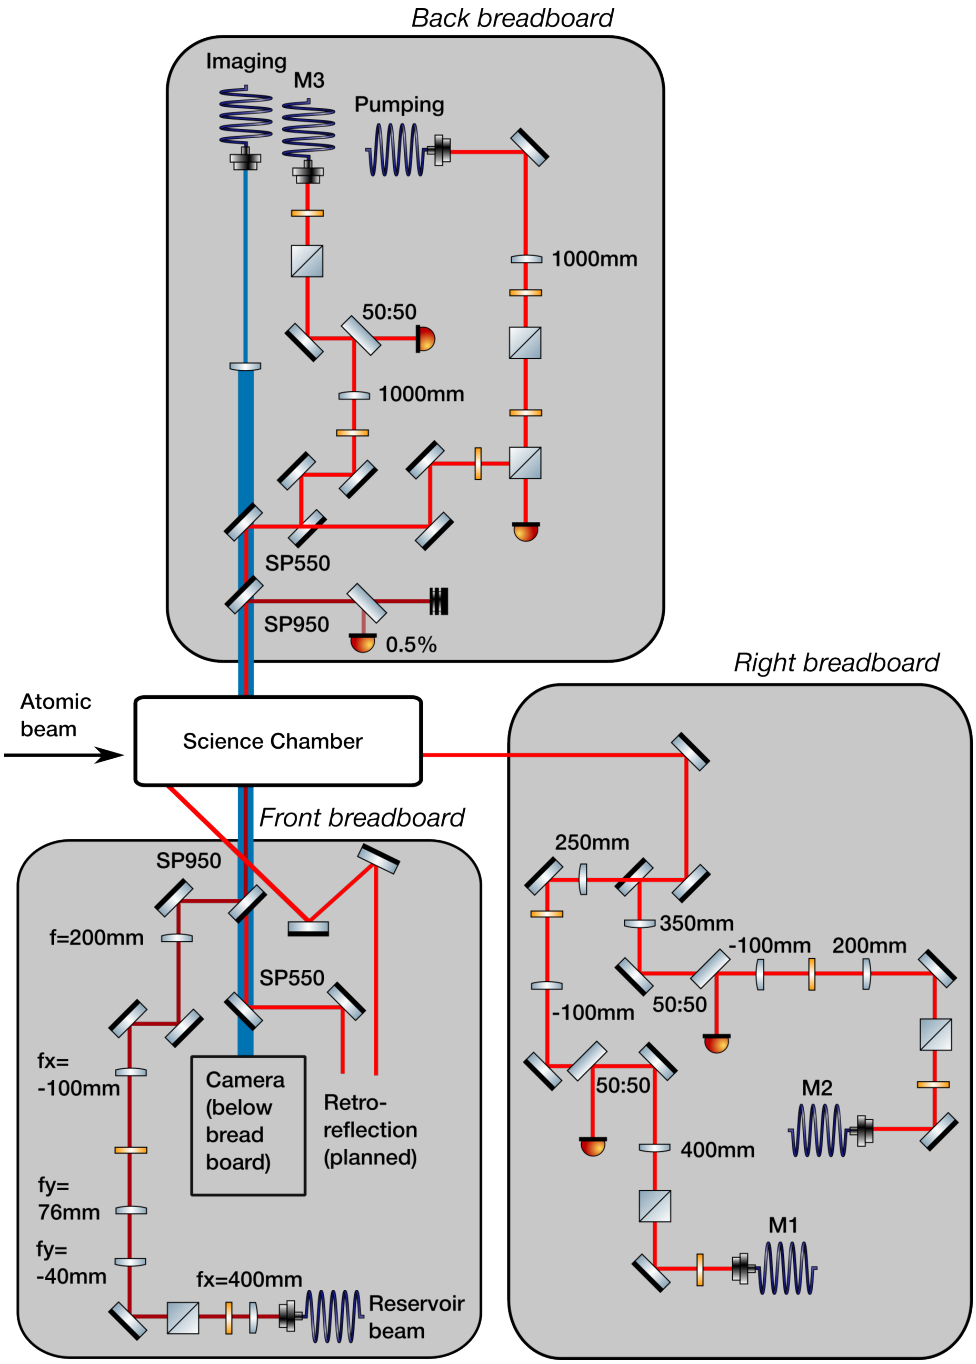


Figure 13: Top view of optical breadboards adjacent to the science chamber (front, back, and right in this top-view), showing the beam paths for the reservoir cooling, pumping and imaging beams. Symbols as defined in the List of Optical Elements (Figure 1).

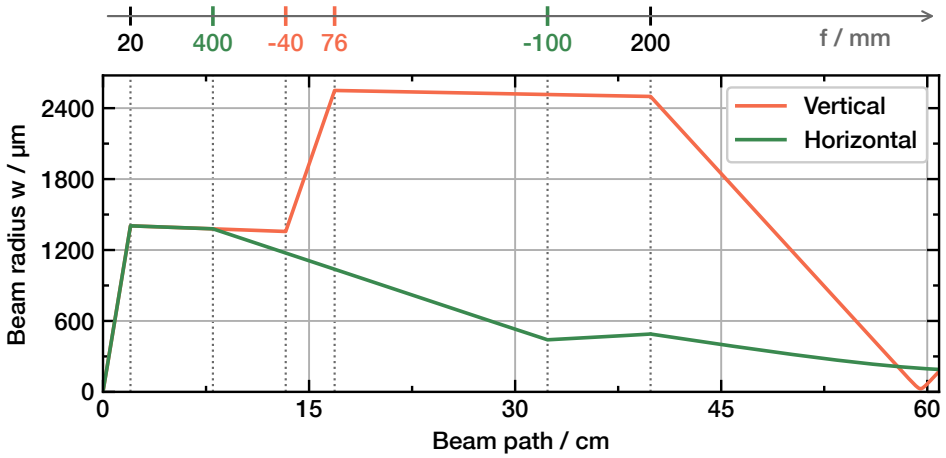


Figure 14: Beam shaping design for the reservoir trap beam. The beam radius along the propagation axis is shown for the horizontal (green) and vertical (orange) directions, with cylindrical lenses acting independently on each axis. The final waists in this model are $w_y = 27\,\mu\text{m}$ and $w_x = 199\,\mu\text{m}$ at the cavity position.

the system matrix. For a thin lens of focal length f and free-space propagation over distance d the transfer matrices are

$$M_f = \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix}, \qquad M_d = \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix}, \tag{4.10}$$

and the output beam parameter is given by $q' = (Aq + B)/(Cq + D)$, where A, B, C, D are the elements of the total system matrix. Lens configurations were designed using this method, implemented in the beam propagation software *reZonator*[1], and verified with measured beam profiles.

The crossed reservoir dipole trap beam is launched from the front breadboard. It requires an elliptical focus, with a tight vertical waist of $w_y \approx 30\,\mu\text{m}$ to achieve sufficient trapping depth, and a larger horizontal waist $w_x \approx 200\,\mu\text{m}$ to ensure good spatial overlap with the cavity mode. The two axes are shaped largely independently using cylindrical lenses, as shown in Figure 14, which shows the evolution of the two beam radii along the beam path. The beam shaping uses more lenses than strictly necessary because the fiber mode-field diameter was not precisely known at the design stage, which required an additional corrective telescope. The beam radii measured at the distance corresponding to the cavity position, $w_x = 192\,\mu\text{m}$ and $w_y = 32\,\mu\text{m}$, agree well with the model values of 199 µm and 27 µm.

[1] http://www.rezonator.orion-project.org/

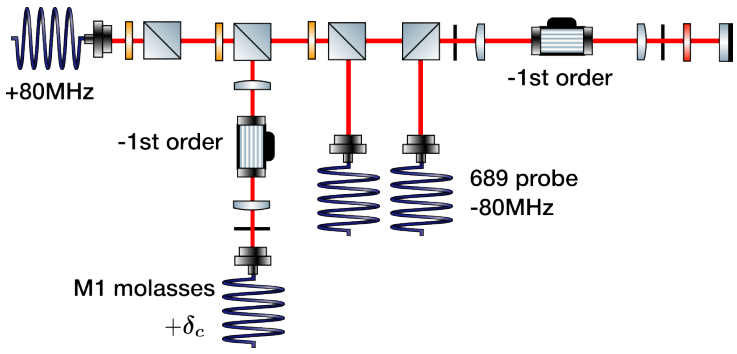


Figure 15: Reduced schematic of the 689 nm distribution and frequency-shifting setup, showing the two output paths to the cavity probe (double-pass) and molasses (single-pass) beams with their respective AOM shifts. Symbols as defined in the List of Optical Elements (Figure 1).

Each beam path includes a photodiode for power monitoring and if needed, intensity stabilization. Polarization is cleaned after each fiber output using a half-wave plate and polarizing beam splitter, and a waveplate close to the science chamber sets the final linear polarization orientation according to [32].

4.4 Supporting Optical Systems

Two additional optical setups were built to support the science chamber infrastructure: a distribution and frequency-shifting setup for the 689 nm light, and a test setup for retroreflecting beams, such as the molasses beams M2 and M3, and the pumping beam. These are described in the following subsections.

4.4.1 689 nm Light Distribution and Frequency Shifting

The first molasses beam and the cavity probe beam both require 689 nm light near the intercombination resonance, but serve different purposes and have different requirements on tunability. The molasses beam addresses atoms across a range of velocities during cooling, requiring a broadband frequency modulation, while the cavity probe beam addresses the narrow, kHz-linewidth transition and therefore requires precise frequency tunability. Its first use is to characterize the atom-cavity coupling by measuring the normal-mode splitting, scanning the probe across the coupled resonance. The same tunability is later used to probe superradiance.

The injection-locked child laser supplying both beams operates at +80 MHz above the atomic resonance, so both paths require a downward frequency shift. The distribution and frequency-shifting setup is shown in Figure 15, with a total input power of approximately 10 mW, of which each path requires at most 0.5 mW. A third output is reserved for future use.

For the cavity probe, the required tunability is achieved using a double-pass AOM configuration, shifting the light by $-2 \cdot 80\,\mathrm{MHz}$. The double-pass setup keeps the beam spatially aligned during frequency scanning. A separate AOM on the cavity breadboard above the science chamber provides the final shift to bring the light to resonance. For the molasses beams, a single-pass AOM at $-80\,\mathrm{MHz}$ shifts the light from the child laser back to near resonance and applies the broadband frequency modulation described in Section 4.1.

4.4.2 Retroreflection Beam Setups

As mentioned in the previous chapter, several beams must be retroreflected to target the atoms in the cavity from both directions, where the quality of this retroreflection is crucial for the cooling and pumping performance. A test was therefore performed to decide which retroreflection scheme to implement.

The goal of the retroreflection setup is to reproduce the incoming beam as closely as possible, in both propagation axis and beam profile. The simplest configuration is a flat mirror at normal incidence, which reverses the propagation direction but does not compensate for beam divergence, so the retroreflected beam is wider than the incoming beam at the target position. A cat-eye (CE) configuration, in which a lens is placed one focal length in front of a flat mirror, additionally inverts the beam divergence and can therefore reproduce the beam waist at a chosen position. A proposed improvement replaces the flat mirror with a curved mirror of radius of curvature $\mathrm{ROC} \approx 1.51f$, where f is the focal length of the lens, which extends the range of incidence angles over which the retroreflection remains effective [41]. Here, a lens of focal length $f = 35\,\mathrm{mm}$ and a curved mirror with $\mathrm{ROC} = 50\,\mathrm{mm} = 1.43f$ were used. This was the closest option to the theoretical optimum of 1.51 that was practically available with commercial optics.

Three configurations were tested to compare the quality of the retroreflection, shown in Figure 16. A slightly diverging beam is incident on a 50:50 beam splitter where the reflected arm serves as a reference, monitored by a camera, while the transmitted arm passes through the retroreflection setup and is partially reflected back through the beam splitter also onto a camera. Both camera positions are at equal distances from the beam splitter, so that a perfect retroreflector would produce identical beams in both arms. The beam used for this test was at 689 nm, using the pumping beam's beam-shaping optics, and retroreflected at a distance of approximately 30 cm behind the cavity. The requirements for the molasses beams are similar, so this test is representative for those beams as well. Alignment was performed by overlapping the back-reflected beam with the incoming beam. For the cat-eye configurations, the lens-mirror distance was tuned to match the beam radius in the far field at 45 cm. Beam radii were recorded at distances of 10 cm, 32.5 cm, and 45 cm from the beam splitter.

The results are shown in Figure 17(a). Comparing the reference beam size with the retroreflected beam size indicates how well the beam divergence is inverted by each retroreflection

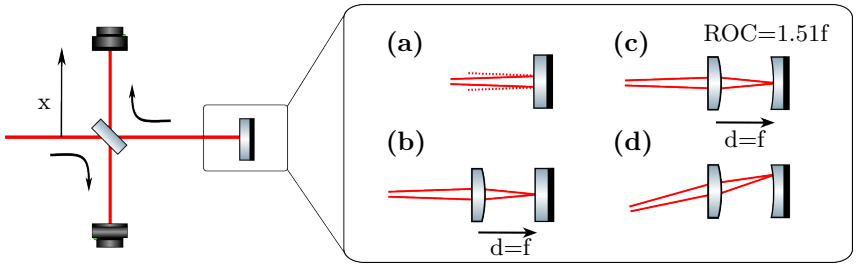


Figure 16: Optical setup used to compare different retroreflection schemes (a-d). A 50:50 beam splitter (BS) separates the incident reference beam from the retroreflected beam. Both beams are recorded at the same propagation distance after the BS using a camera. (a) Single flat mirror. (b) Conventional cat-eye (CE) configuration consisting of a lens and a flat mirror. (c) Cat-eye configuration with a curved mirror (optimized ROC). (d) Beam propagation inside the curved-mirror CE for angled incidence.

scheme. As expected, the flat mirror produces a retroreflected beam that is systematically larger than the reference over the full propagation range, since the divergence is not compensated. Both cat-eye configurations match the reference beam size at 45 cm, where alignment was optimized, and show small deviations at shorter distances, likely due to measurement imprecision at the alignment distance. Beam radius measurements from the camera images carry some uncertainty due to strong interference fringes on the sensor and potentially angled camera placement. The values were fitted using Equation (2.9) with the known beam waist $w_0 = 327\,\mu\text{m}$. Within this measurement precision, the conventional and curved-mirror CE setups perform comparably for this slightly diverging beam.

The effect of incidence angle and beam divergence were investigated as additional checks. Tilting the curved-mirror CE setup with respect to the incoming beam showed that beam size matching is maintained over a wider range of angles than would be possible with the flat-mirror CE, consistent with the design idea of [41]. At an incidence angle of 14°, shown in Figure 17(a), the curved-mirror CE remains limited mainly by the geometric displacement of the beam, which must still strike the mirror. Because the beam meets the curved surface at a locally perpendicular point, it is still reflected back in the same direction. The conventional (flat-mirror) CE, by contrast, no longer reflects the beam back in the same direction at this angle. For a more strongly diverging beam, produced by inserting an additional 200 mm lens, both CE setups continue to outperform the flat mirror, as shown in Figure 17(b). The curved-mirror CE shows a somewhat larger deviation from the reference in this case, likely due to spherical aberrations which become more significant at larger beam size and were neglected in the paraxial model of [41].

Camera images were additionally analyzed for aberrations. The ellipticity increases somewhat for the tilted curved-mirror CE but remains small for normal incidence. The sum of

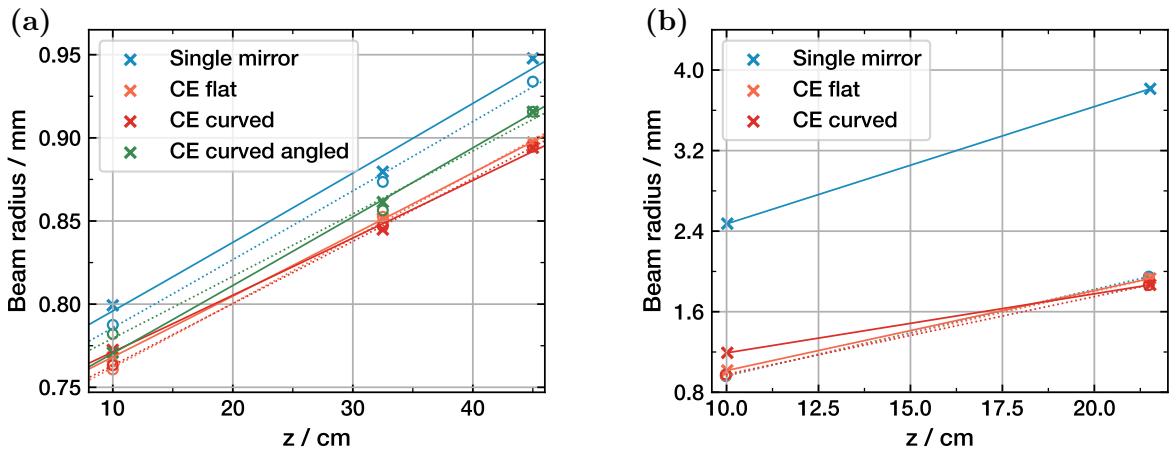


Figure 17: Measured beam radii of the directly reflected reference beam (open markers) and the retroreflected beams (cross markers) for different retroreflection configurations. (a) Retroreflection of a slightly diverging beam using a flat mirror, a conventional cat-eye (CE) setup, and a curved-mirror CE setup where measurements are shown for normal incidence and for an incidence angle of 14° together with Gaussian fits. Solid lines correspond to the retroreflected beams and dotted lines correspond with directly reflected reference beams. (b) Measured beam radii for a strongly diverging beam in the corresponding retroreflection configurations, lines connecting the measurement points.

fit residuals increases slightly for both CE configurations compared to the flat mirror, and marginally more for the tilted case, indicating a small but measurable growth in wavefront distortion. A full aberration analysis is limited by interference fringes on the camera sensor.

Overall, these measurements confirm that a cat-eye configuration is essential for retroreflecting the molasses and pumping beams, and that the conventional CE setup is sufficient for the low-divergence beams that can be freely aligned in this experiment.

5 | LASER INTENSITY STABILIZATION

As explained in the introduction, a big advantage of the superradiant laser is that the output frequency is set by the atoms rather than by the stability of an external cavity. This can remove cavity-length fluctuations as the limiting factor, which now shifts the focus to other noise sources that perturb the atoms themselves.

An important example is the light shift induced by the laser beams controlling the motional and internal states of the atoms. Unless a beam operates close to a magic wavelength, where the shifts of both states cancel, its AC Stark shift (Section 2.1.2) changes the transition frequency. For the pumping beam, simulations of this experiment predict only a minor influence of the pumping-induced light shift on the superradiant output [32]. More critical are the 1070 nm dipole guide and reservoir beams, which shift the cooling transition by position-dependent amounts of up to the MHz level, far exceeding its linewidth [32]. The cooling frequencies at the reservoir can be adapted to these shifts, but intensity fluctuations displace them (already at the percent level by more than the transition linewidth) and thereby directly perturb the cooling and loading of the atoms.

As discussed in Section 2.3, intensity fluctuations on cooling and trapping light can translate into heating of the atomic sample, limiting the achievable temperatures and atom numbers. This is particularly relevant for the reservoir beam, which holds the atoms as they are cooled before being transported through the cavity. At the level of the MOT beams, intensity fluctuations introduce noise on the position of the trap center [42]. This is especially relevant for the vertical MOT beams, where any imbalance shifts the trap center upward, as the stronger beam from below must compensate gravity. These position fluctuations lead to variation in the number of trapped atoms and reduce the overlap with the dipole guide, affecting the loading efficiency.

All of these effects ultimately contribute to the same problem: they cause fluctuations in the number of atoms delivered to the emission zone. Simulations predict that such atom number fluctuations are the main mechanism broadening the linewidth of the superradiant output [32]. Stable laser intensity throughout the preparation and transport chain is therefore a requirement for stable superradiant emission, and active intensity stabilization is essential.

This chapter presents the implementation and characterization of the intensity stabilization systems intended for this experiment. I first introduce the test setup and working principle of the feedback control loops. I then characterize and compare the different PID systems implemented, evaluating step responses and noise spectra using parameters representative of the pumping beam and the frequency ranges relevant to this experiment. However, much of the underlying characterization is more general and can serve as a reference for other PID-based stabilization systems.

5.1 Proportional-Integral-Derivative Controller and Experimental Setup

A proportional-integral-derivative (PID) controller is a fundamental system in feedback regulation of dynamical systems. Its goal is to make a system's process value (PV) follow a desired setpoint (SP) by adjusting a control value (CV) that steers the system dynamics in the required way. An intuitive example is a thermostat which measures the current temperature (the PV), compares it to the desired temperature (the SP), and adjusts the heating output (the CV) to maintain the target temperature. For stabilizing laser intensity, the process value is the measured laser power on a photodiode. An acousto-optic modulator (AOM) provides control over the intensity by diffracting the light into different orders depending on the applied RF power, which is set via the control voltage.

The PID algorithm determines the control value output by combining three contributions addressing different aspects of the error between setpoint and process value $e(t) = \mathrm{SP}(t) - \mathrm{PV}(t)$.

- The proportional term, $K_P\, e(t)$, reacts to the current error, producing a correction proportional to its size.
- The integral term, $K_I \int_0^t e(\tau)\, d\tau$, accumulates the error over time, ensuring that any small persistent offset, or "proportional droop", is gradually eliminated.
- The derivative term, $K_D \frac{de(t)}{dt}$, reacts to the rate of change of the error, which can anticipate and reduce overshoot of the process value after sudden changes.

By tuning the gains K_P, K_I, and K_D, one can adjust how much each term contributes, balancing speed, accuracy, and stability. Typically, proportional and integral gains are essential for achieving a fast and accurate feedback response, while the derivative gain can destabilize the system and is not always used. A helpful introduction to the topic of feedback control is given in [43].

The setup used to test the PID for intensity stabilization is shown in Figure 18. Laser light at a wavelength of 689 nm is polarization-filtered and focused through an AOM to minimize the acoustic wave transit time and achieve a fast switching response. The first-order diffracted beam is coupled into a polarization-maintaining fiber, while the zeroth order is blocked. This choice is made because the intensity in the first order can be varied continuously between approximately 0 % and 70 % of the total input intensity, whereas the zeroth order can not be fully extinguished. The fiber itself introduces intensity noise: imperfect polarization maintaining and beam pointing instability cause the polarization state to drift, and this drift is converted into power fluctuations by the polarization-cleaning optics after the fiber. The light is split between the experimental signal, detected here with a biased photodiode (Thorlabs DET10A2, 1 ns rise time), and the feedback signal, detected with an amplified photodiode (Thorlabs PDA10A2,

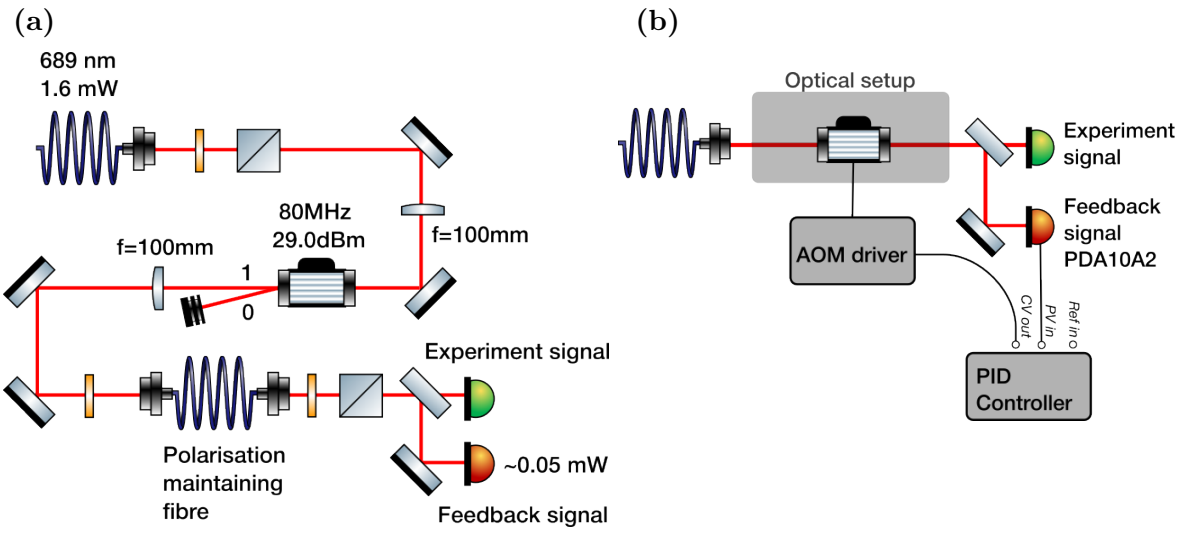


Figure 18: Schematic of the experimental setup used for intensity stabilization and its control structure. The optical setup (a) prepares the beam and modulates the first-order AOM output intensity, while the control structure (b) illustrates the PID feedback loop that regulates the AOM driver based on the photodetector signal. Symbols as defined in the List of Optical Elements (Figure 1).

150 MHz bandwidth). The latter is sent in a feedback loop to a PID controller that drives the AOM's radiofrequency source, thereby regulating the optical power in the first order.

This method inevitably introduces some power loss, the first-order diffraction efficiency of the AOM, the fiber coupling efficiency, and splitting off the feedback signal all place limits on the available power. In practice, however, the intensity stabilization can be combined with frequency shifting and fiber coupling to the experiment table that would be required regardless. In our case the power requirements for cooling and pumping are low and reservoir beam power is not a limiting factor, so this is not a concern here.

For driving the AOM, two options are available: a commercial AOM driver, and a DDS combined with a separate RF amplifier. The first option is the commercial device (DIM-3000, from Time Base), which drives the AOM directly without an external RF input signal. Its amplitude modulation input allows adjustment of the output power using a control voltage between −1 V and 1 V. A control voltage of 0 V corresponds to unmodulated operation at the set RF power, while 1 V results in power gain and −1 V in full attenuation. Using an amplitude modulation offset in arbitrary units, the maximum output power can be shifted across input control voltages. The second option is a rack-compatible RF amplifier, constructed in the Technology Center of the University of Amsterdam. This amplifier requires an external RF

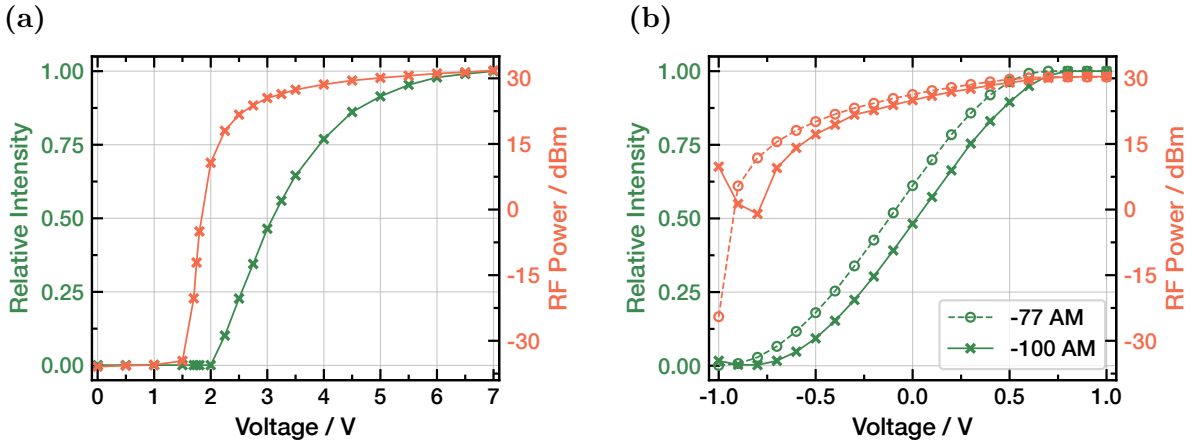


Figure 19: RF driver transfer functions: RF power and intensity in the first AOM diffraction order as functions of the applied control voltage for (a) the RF amplifier with a fixed RF input power of −11.1 dBm and (b) the DIM-3000 at a set RF power of 29 dBm, where two different amplitude modulation (AM) settings are shown.

input set to the desired AOM frequency and a control voltage between 0 and 7 V to generate an amplified RF output.

Figure 19 shows a characterization of both AOM-drivers including the output power for the full range of input control voltages and the corresponding measured laser intensity in the first order after the fiber. Both, the RF tuning via voltage input, and the AOM itself introduce a non-linearity which can limit the PID outcomes [43]. Notably, the RF amplifier barely gives any intensity in the first order, until 2 V are reached, which limits the usable voltage range. For the DIM-3000 measurement, two different amplitude modulation offsets were analyzed, both for a power setting of 29 dBm. The optimal setting with the biggest output range for the RF power was found to be −77. In the following PID characterization, the offset was set to −100 , which shifts the intensity curve slightly down and gives an optimal behavior at the used intensity values. This introduces a non-monotonous behavior for very low voltages which is however barely measurable in the optical domain.

We further have two choices for the PID controller, both constructed in the Technology Center of the University of Amsterdam. The first is a Red Pitaya[2]-based digital PID which can be controlled and tuned with the software *pyrpl.exe*[3]. Second, an analog PID on which the PID settings can be tuned using physical controls. In the analog PID, the error signal is determined entirely with electronic components such as resistors, capacitors and operational amplifiers. In the digital PID, the input signal is first digitized, the PID calculation is performed digitally, and the resulting output is converted back to an analog signal. Given the two PID modules and

[2] https://redpitaya.com
[3] https://pyrpl.readthedocs.io

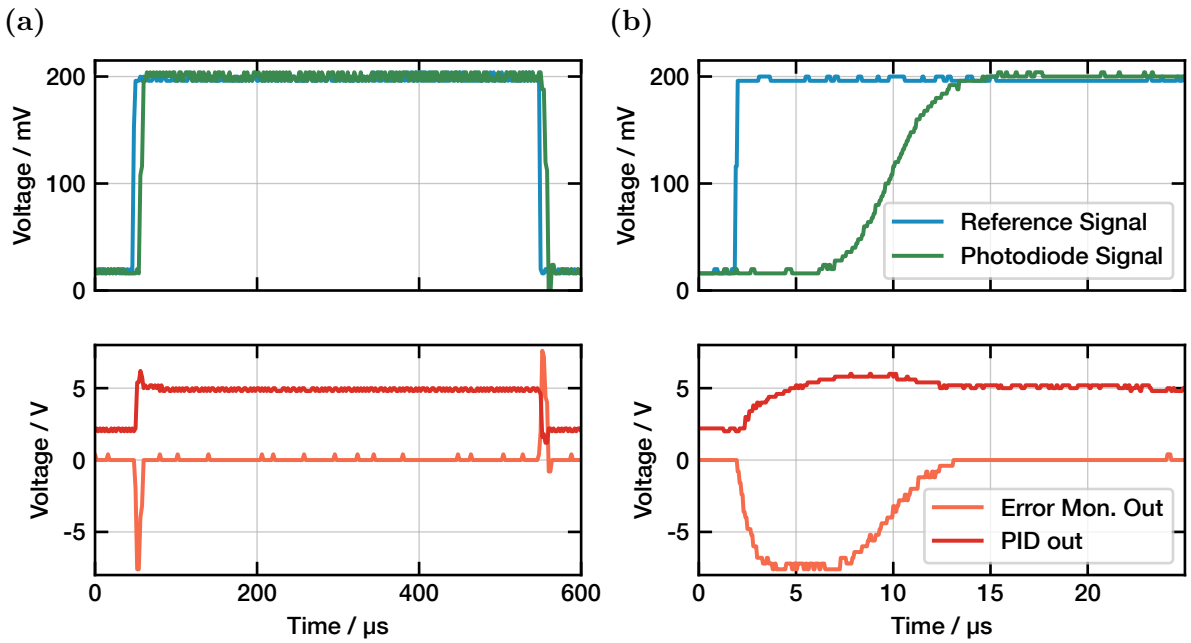


Figure 20: Example measurement illustrating the PID working principle using a step-function setpoint (RF amplifier with analog PID). The upper plots show the setpoint reference signal (blue) and the photodiode signal (green), while the lower plots display the internal PID operation: the error signal (orange) and PID output (red). (a) Full step function, (b) zoomed view of the rising edge.

two possible AOM drivers, four different combinations for the feedback loop can be tested. The output voltage range of the PID system must be internally adjusted to match the modulation range of the AOM driver used.

5.2 Step Response Analysis

The working principle of the PID system is illustrated in Figure 20. The external setpoint (blue) and the photodiode signal (green) go to the PID controller which calculates the error (orange) as their difference. As the setpoint voltage jumps up, the error signal initially drops to a large negative value, causing the PID output to rise rapidly. If the PID gains are well tuned, both the output and the photodiode signal stabilize quickly after this adjustment period, continuously comparing the measured signal with the reference. The falling edge of the step function produces the opposite response, exhibiting a slight overshoot that was resolved in a newer version of the RF amplifier. On the right-hand side, a zoomed-in view of the rising edge illustrates the dynamics more clearly.

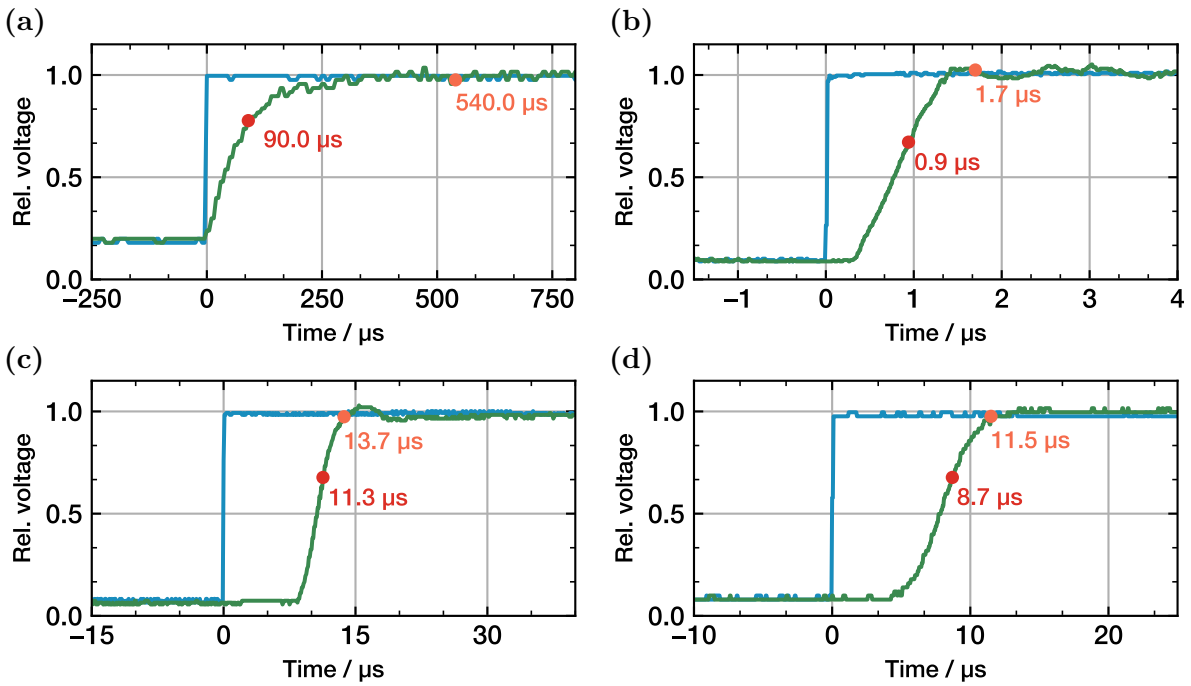


Figure 21: Step responses of the four PID loop systems showing relative voltage: blue lines indicate the setpoint, green lines the photodiode signal. $1 - 1/e = 63.2\,\%$ rise time (red) and long-term stability $\pm 3\,\%$ (orange) are indicated. Panels correspond to: (a) Digital PID, DIM (b) Analog PID, DIM (c) Digital PID, RF amplifier (d) Analog PID, RF amplifier.

The tuning of the PID parameters, although implemented differently in the two PID modules, is done with the same approach. Following the general idea of the Ziegler-Nichols method [44], the optimal parameters are determined. First, the integral and derivative gains are set to zero, and the proportional gain is gradually increased until the photodiode signal exhibits large oscillations. This value is then slightly reduced, and the integral gain is increased until the step response is as fast as possible without a significant overshoot. Derivative gain was tested on all systems but did not provide reliable improvement.

For many laboratory applications, it is important that the setpoint is tracked quickly without adding oscillations to the system. Therefore, the system's response to a step-function input is tested, which makes the speed of the control loop easy to evaluate directly. The setpoint is stepped by around 250 mV, to below the photodiode's maximum signal, leaving the PID room to adjust the output in both directions. Later, in the noise spectrum analysis, we will consider a constant setpoint, where the photodiode signal fluctuates on a much smaller scale. The underlying working principle, however, is the same.

Examining the step-function responses in Figure 21 reveals clear differences between the four combinations. Rise and long-term settling times (within $\pm 3\,\%$) were extracted for each

combination. For reference, the bare response time of the two AOM drivers without any PID loop sets a lower bound on how fast a closed loop using that driver could in principle respond. The RF amplifier alone has a $1/e$ rise time of 1.4 µs and settles within ±3 % at 3.5 µs, while the DIM-3000 alone is considerably faster, rising in 0.16 µs and settling at 0.19 µs. Since both drivers address the same AOM, contributions from signal propagation and the acoustic transit time in the crystal are common to both measurements and are not separated here. They are negligible compared to the RF amplifier's response, but may contribute at the level of the DIM-3000.

To compare the rise times to bandwidths, typically defined as the frequency at which the input and output signals develop a 45° phase lag, we convert the former using $f_{3\mathrm{dB}} = 1/(2\pi\tau)$, valid for a system with a single dominant time constant [43].

The DIM RF-source with the digital PID is significantly slower than the other three combinations due to insufficient integral gain. This appears to be a tuning limitation of the RF curve of this particular combination rather than a fundamental constraint of the digital PID. Given its response time of 90 µs, the noise spectrum was not measured in the following. The DIM RF-source with the analog PID responds fastest of the four on both metrics, with a rise time of 0.9 µs, due to much higher achievable integral gain. The RF amplifier with both digital and analog PIDs shows intermediate performance (11.3 and 8.7 µs rise times, respectively). In the digital case, the integral gain is at its limit again, leading to a slight overshoot. Comparing these rise times to the bare response of the drivers shows that the PID loop itself is the limiting factor in both cases: the RF-amplifier combinations stay close to the bare driver's rise time, while the DIM combinations remain considerably slower than the bare DIM driver, which on its own is faster. Since both PID types perform well with the RF amplifier but the digital PID performs poorly specifically with the DIM, the issue is unlikely to lie with either component individually. An explanation is a suboptimal operating point on the DIM's RF-power transfer curve in this configuration, suggesting that further tuning of the power and offset settings could improve the digital PID's performance with this driver.

5.3 Noise Spectra Analysis

Now using a constant setpoint, we investigate how the noise spectrum changes when PID control loops are implemented. Noise spectra were recorded using the PNA1 Thorlabs intensity noise analyzer, suitable for low frequency noise, and averaged over 20 measurements. Figure 22 compares the relative intensity noise (RIN) for the three operating PID and driver combinations identified in the previous section against the open-loop (no PID) case. Data from the noise analyzer (for lower frequencies) and a separate spectrum analyzer (for higher frequencies) are combined to cover the full frequency range, with reduced confidence regions shown with partial transparency.

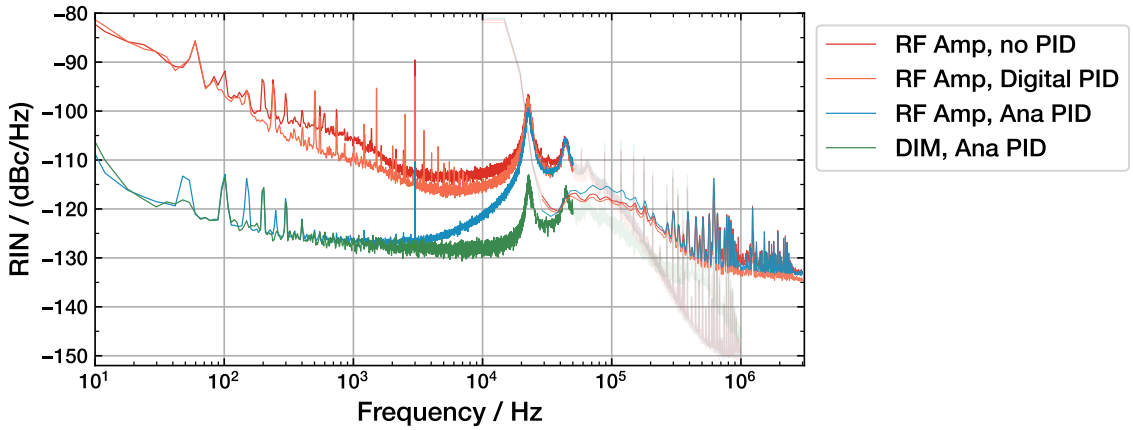


Figure 22: RIN noise spectra of the three operating PID/driver combinations and the open-loop (no PID) case. Data from the PNA1 intensity noise analyzer (lower frequencies) and a spectrum analyzer (higher frequencies) are combined to cover the full range, with reduced-confidence regions shown with partial transparency.

Initial attempts to further optimize the PID gains based on the noise spectrum were not successful. In particular, the hypothesis that a smaller proportional gain would improve the noise performance could not be confirmed. Best gain tuning for noise spectrum was ultimately based on the step response in all cases. Comparing the three combinations, the analog PID provides the strongest low-frequency noise suppression, reducing the noise by around 25 dB across a broad frequency range, relative to the open-loop case. With the RF amplifier, this analog PID suppression begins to slowly degrade above roughly 3 kHz, below the combination's calculated bandwidth of $(2\pi \cdot 8.7\,\mu s)^{-1} = 18.3\,kHz$. With the DIM driver, the suppression persists to higher frequencies, consistent with its larger bandwidth ($(2\pi \cdot 0.9\,\mu s)^{-1} = 177\,kHz$). The digital PID with the RF amplifier shows only a modest noise reduction in the 0.1 to 10 kHz range.

The limited noise reduction of the digital PID is likely due to additional electronic noise. The conversion between analog and digital signals can lead to so-called spurious effects from the sampling. However, these would be expected to appear as narrow peaks rather than raise the noise floor across the spectrum. Instead, the broadband nature of the noise suggests that the Red Pitaya platform, not optimized for low-noise operation in this frequency range, is the limiting factor for the digital PID's performance here.

At higher frequencies, in the hundreds of kHz range, the noise analyzer's photodiode (Thorlabs DET10A2) becomes unreliable, since it is battery-biased and relies on a passive RC low-pass filter rather than a transimpedance amplifier [45]. This lacks an actively stabilized bias and limits the detector's usable bandwidth. The cable-powered photodiode Thorlabs PDA8A2 (50 MHz bandwidth) with a spectrum analyzer was used instead to verify this regime. The relative in-

tensity noise was calculated from the spectrum analyzer data as follows. The carrier power P_c was determined from the measured DC photodiode voltage V_{DC} across the $R = 50\,\Omega$ input impedance, $P_c = V_{DC}^2/R$, and converted to dBm. The RIN, in dBc/Hz, was then obtained from the measured noise power P in dBm, as

$$\mathrm{RIN} = P_{\mathrm{dBm}} - P_{c,\mathrm{dBm}} - 10\log_{10}(\mathrm{ENBW}), \tag{5.1}$$

where the equivalent noise bandwidth ENBW was taken equal to the analyzer's resolution bandwidth of 10 kHz. This confirms that the noise does not continue to decrease at high frequencies, as the noise analyzer measurement alone might suggest, but instead levels off. The two measurement methods do not agree in absolute noise level, but are consistent enough for a qualitative comparison.

All measured noise spectra reflect several physically distinct noise sources, dominant in different frequency regimes. At low frequencies, up to a few hundred Hz, mechanical and acoustic vibrations of the optical setup and surrounding lab environment dominate [46]. Above this range, electronic noise from the control and detection electronics and from other lab equipment, together with the laser's intrinsic intensity noise (e.g. relaxation oscillations), becomes relevant. The peak at 22 kHz visible in the spectrum can be identified as the servo bump of the laser lock, since it is also present in the open-loop (no PID) measurement and is therefore not introduced by the intensity-stabilization loop itself. At high frequencies, like the MHz range, the noise floor is expected to converge toward the fundamental shot-noise limit set by the quantized arrival of photons at the detector [47]. These contributions could be separated better after a comparison with the noise spectrum of an incoherent white-light source at the same optical power, but no such source was available.

A number of practical factors were found to affect the measured noise floor. Ground loops in the cabling introduced additional low-frequency noise, visible in Figure 23(a) as peaks at 50 Hz and its (particularly odd) harmonics. Such loops arise when a path is grounded at more than one point with a potential difference between them. Here, the AOM driver and the setpoint source from the DAC, powered from a separate supply in the rack, created an unwanted current loop. Using the PID's internal setpoint, rather than an externally supplied one, reduced this issue by removing one of the competing ground connections. In addition, all relevant devices were connected to the same power outlet, eliminating further potential differences between their grounds. The magnitude of these peaks illustrates the impact ground loops can have on the noise floor, and highlights the importance of checking external setpoint signals for this issue when implementing a similar stabilization system.

Additionally, the influence of the intensity on the feedback photodiode was tested. The intensity on the noise-analyzer photodiode was kept constant, while the optical intensity in the feedback loop was varied between 0.05 mW (resulting in a 180 mV photodiode signal) and

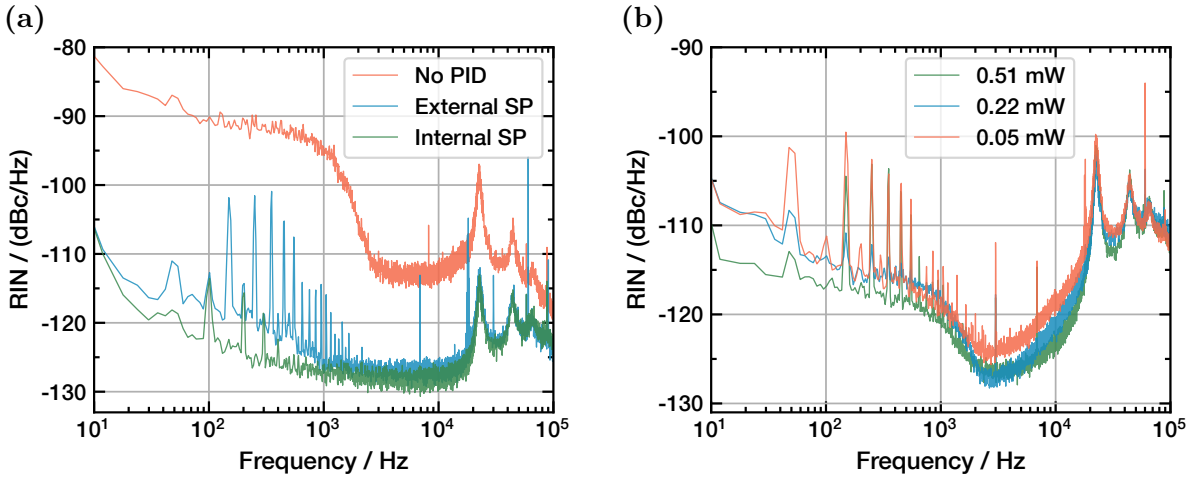


Figure 23: (a) RIN noise spectrum showing the difference between pronounced ground loop peaks of the cabling from an external setpoint, the reduced peaks for internal setpoint of the PID and for reference the open loop noise spectrum (Analog PID with DIM driver). (b) RIN noise spectra for varying optical power on the feedback photodiode (labeled in legend), with the intensity on the noise-analyzer photodiode held constant. Note the different vertical scale compared to Figure 22.

0.51 mW (resulting in a 2.0 V photodiode signal), with results shown in Figure 23(b). At lower optical power on the feedback photodiode, the signal to noise ratio of the photodiode decreases, degrading the intensity stability. To check whether this is limited by the fundamental shot-noise floor, we compare it to the expected shot-noise-limited RIN, $\mathrm{RIN}_{\mathrm{shot}} = 2hc/(\lambda P)$, where hc/λ is the photon energy at $\lambda = 689\,\mathrm{nm}$ and P is the optical power on the photodiode [47]. This predicts shot noise floors of $-139\,\mathrm{dBc/Hz}$ at 0.05 mW and $-149\,\mathrm{dBc/Hz}$ at 0.51 mW, a 10 dB reduction for a tenfold increase in power. The measured noise floor lies above this prediction at both powers and decreases by only a few dB between them. This shows that the stabilization is not shot-noise-limited in this regime. The observed power dependence instead reflects electronic noise in the photodiode and its readout, which becomes more significant relative to the photocurrent at lower optical powers.

Long-term drift and stability were also assessed by recording the normalized intensity over 5.5 hours on a power meter, shown in Figure 24. Both the analog and digital PIDs with the RF amplifier show reduced long-term variance compared to the open-loop case, though some residual drift and few short dips remain. These dips and potentially the long term drift are likely due to the limited reliability and stability of the Thorlabs power meter used for this measurement, rather than to the stabilization loop itself.

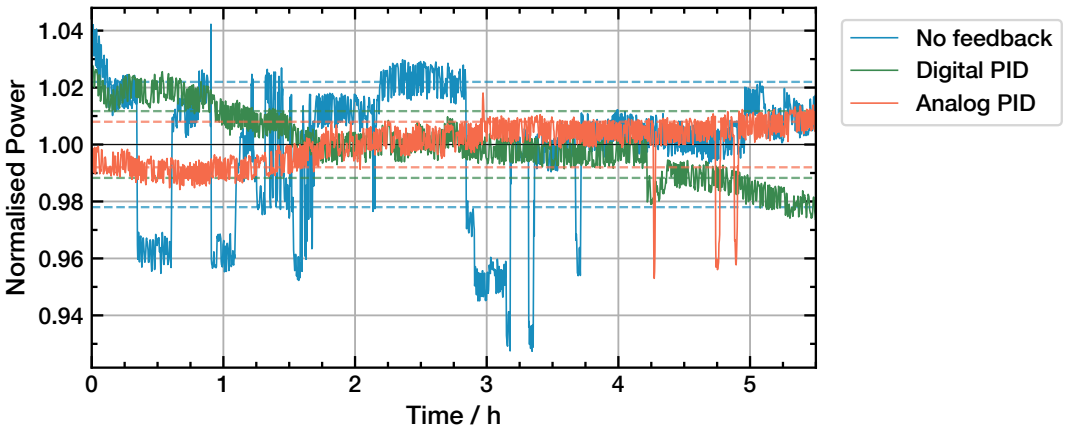


Figure 24: Normalized laser intensity over a 5.5 hour period for the open-loop, analog PID, and digital PID cases (RF amplifier driver), illustrating long-term drift and stability and indicating the variances in the power distribution with dashed lines.

5.4 Sample and Hold Function

The analog PID includes a sample-and-hold (S&H) function via an additional TTL input that can be enabled by switching a pull-down resistor in the PID electronics. When activated, the feedback loop is paused and the PID output held constant at its last value, while the AOM continues to be driven at the corresponding RF power. This function was tested as a potential tool for situations in the experimental sequence where the stabilization loop should not react to changes in the photodiode signal.

Since the PID output goes through the S&H chip, whenever this function is enabled, even outside of hold periods, its presence could add a delay to the loop and decrease the PID bandwidth. To check this, the step response with the S&H chip in the signal path was compared to the response without it. The rise time was unaffected, while the fall time increased slightly, from 3.8 µs to 5.6 µs. Figure 25(a) compares the noise spectrum with and without the S&H function enabled, showing its effect on the intensity noise. The RIN increases slightly starting around 1 kHz, and more substantially above roughly 100 kHz. Since this lies well above the trap-frequency band, this increase is not a significant concern here.

The working principle of the hold function is illustrated in Figure 25(b). While the S&H input (blue) is low, the loop is paused. During this, the RF drive to the AOM (green) is held constant and no longer responds to the photodiode signal (orange). When the loop is restarted, an overshoot of around 50 mV for 50 µs appears before the system recovers, which should be kept in mind for applications where this could affect the atoms. The inset of Figure 25(b) shows that the photodiode signal rises slowly over the first second after the hold is released by a small extend. This very slow dynamics is most likely due to thermal heating of the AOM during the

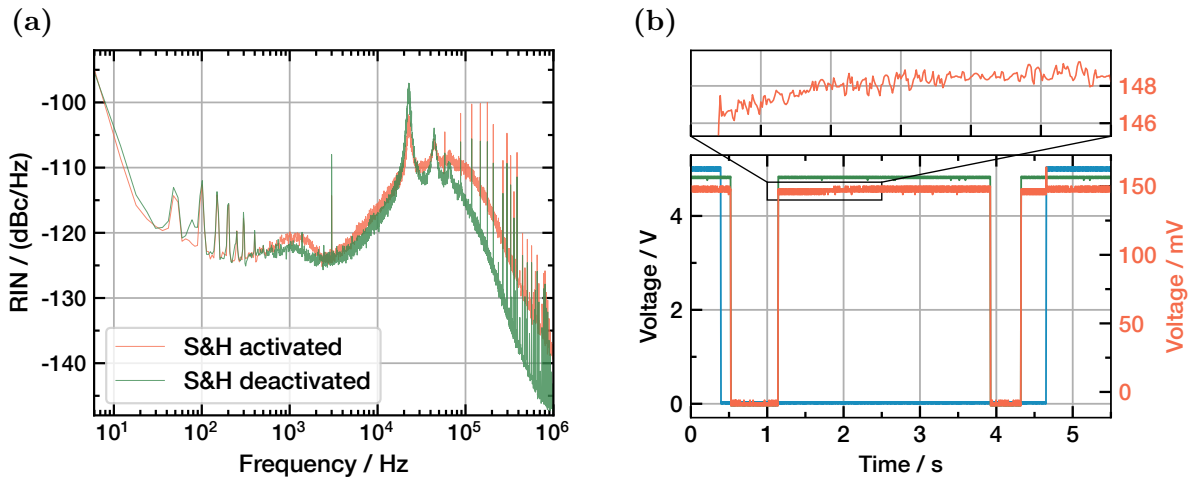


Figure 25: Sample-and-hold (S&H) function in PID control and its effect on intensity noise. (a) Noise spectrum with and without the S&H function activated, showing its influence on the intensity noise. (b) Measured voltage levels showing PID behaviour during hold mode: a low S&H input (blue) pauses the loop, the RF drive (green) controls the AOM directly, and the PD signal (orange). The inset shows the slow rise of the PD signal over the first second after hold.

hold period, which would gradually shift its diffraction efficiency once RF drive resumes. An alternative approach to avoid this would be switching the AOM between two RF frequencies, of which only one couples into the fiber, rather than switching the RF amplitude on and off, since the AOM would remain continuously driven. This was not tested here, as the available fast switch did not support this option. In the future, an improved switch will be implemented in the RF amplifier to enable this fast switching.

Concluding Design Choices for this Experiment

Based on the characterization presented in this chapter, the following implementations for the experiment are suggested. Among the two AOM drivers, the RF amplifier is preferred over the DIM-3000 for practical reasons. It is rack-compatible, allowing multiple units to be installed and centrally controlled as needed for the several beams requiring stabilization. Between the two PID controllers, the analog PID performed considerably better in this characterization, in particular in the noise analysis. A recent updated version of this module additionally includes an audible warning signal that is triggered if the deviation from the setpoint exceeds a certain threshold (clip detector), indicating that the loop has lost its setpoint. The only minor drawback of the analog PID is that its gain settings are tuned manually via switches and screws and cannot be reproduced precisely. This has not proven to be a significant practical issue as the readjustment of PID gains was straightforward. The digital PID does not have a big effect on intensity noise

suppression in this setup, though it may remain relevant for other applications, potentially for higher frequency ranges. Identifying the specific source of its larger noise, for example by directly testing the noise in the Red Pitaya output was beyond the scope of this work but would be a useful direction for future improvement.

For the reservoir beam, an AOM is already installed in the beam path for the purpose of intensity stabilization, and the control system characterized in this chapter can be implemented directly. For the dipole guide, applying the same scheme is likely not possible. This high-power beam cannot be guided through an AOM, which is the component used to control intensity and realize feedback here. For the pump beam, stabilization will be installed once the frequency scheme for this beam is set up. For the molasses beams, active stabilization is currently not planned, but could be added if passive intensity stability proves insufficient, since these beams already have an AOM in the beam preparation stage and photodiodes at the beam shaping stage available for monitoring. For instance fluctuating light shifts or scattering at the reservoir position could turn out to complicate the cooling and transport scheme. Finally, the sample-and-hold function gives a practical tool for situations in the experimental sequence where the feedback loop should be paused without switching off the AOM drive entirely.

The intensity stabilization scheme presented in this chapter is not limited to the beams discussed in this thesis and could likewise be applied to other beams in the experiment. This includes the 813 nm moving optical lattice beams, where the cavity transmission could serve as the feedback signal and AOMs are already installed in the existing cavity setup, as well as the 689 nm light used to seed superradiance mentioned in Section 4.4.1, which also already includes an AOM in the beam path from the splitting setup.

Overall, the noise suppression demonstrated in this chapter, particularly the strong low-frequency suppression achieved with the analog PID, directly addresses the stability requirements motivated at the start of this chapter. By limiting heating and stabilizing the atom number for superradiant emission, this represents a promising step toward the level of precise control this experiment requires.

6 | CONCLUSION AND OUTLOOK

This thesis presents work towards the atom source needed for a continuous superradiant laser on the mHz-linewidth clock transition in strontium. Such a source must be simultaneously cold enough to suppress Doppler broadening of the narrow transition, dense enough to efficiently generate superradiance, and delivered continuously without large fluctuations. Starting from the theoretical framework of laser cooling and optical dipole potentials, this thesis explains and characterizes the stages that prepare and transport atoms between a MOT chamber and a separate science chamber housing the superradiance cavity. A blue MOT and a dual-position red MOT were optimized and characterized, providing a continuous source of atoms at temperatures of 8.0 µK, with a saturation atom number of 3.7×10^8 and a loading rate of 1.5×10^8 atoms/s in the broadband red MOT stage. Lower temperatures can be reached in the spatially separated single-frequency component, the demonstrated 3.3 µK can be further reduced to the sub-µK regime [33]. To cover the 20 cm distance to the science chamber, a high-power optical dipole guide was characterized theoretically and experimentally prepared. The beam shaping of the reservoir trap and molasses beams needed to recapture and further cool atoms at the science chamber was set up, along with a frequency shifting system. Testing of retroreflection setups showed that the more advanced cat-eye configuration was not necessary for this application, the conventional cat-eye configuration reaching equally good results. Since intensity fluctuations on this light limit both the achievable atom temperature and the stability of potential superradiant emission, an active intensity stabilization system was set up and tested. Comparing a custom RF amplifier against a commercial DIM-3000 AOM driver from Time Base, and a digital against an analog PID controller for each, the analog PID was found to give the strongest performance, suppressing relative intensity noise by more than 25 dB at the frequencies most relevant to trap heating, while the RF amplifier was preferred as the AOM driver for practical reasons. Together, these results establish a basis for the transport chain from initial cooling through to the entrance of the science chamber that the remaining steps toward continuous superradiant lasing will build on.

Extending the Stabilization System. The stabilization scheme from Chapter 5 can be applied to the reservoir trap and to the pumping beam once its frequency scheme is developed. Intensity stabilization can also be extended to the 813 nm moving-lattice beams, using cavity transmission as the feedback signal, and to the 689 nm beam used to probe superradiance, both of which already have AOMs installed. Other open points that could be addressed include directly analyzing the digital PID's broadband electronic noise, and testing an AOM switching scheme based on two different frequencies, which the improved RF amplifier switch makes possible.

Completing the transport chain. With loading into the dipole guide demonstrated, the next step is to set up the Bloch accelerator beams to push the atoms along the guide towards the cavity. Additionally, the loading efficiency of currently around 2 % could be increased by further optimizing the single-frequency MOT for loading into the guide. The focus of the guide beam can then be shifted to the midpoint between MOT and cavity, with the aim of detecting atoms at the cavity position using the existing absorption imaging system. Once atoms are detected there, the molasses beams at the science chamber can be optimized on the atoms in alignment and frequency structure to reach the best cooling and capturing performance. Possible modifications at this stage include adjusting the waist of the guide beam to reduce the axial potential gradient, and verifying the beam profile quality of the high-power beam at the powers used. Should the beam quality prove insufficient, a solution would be spatial mode filtering with a pinhole in the Fourier plane of a telescope, with the telescope lenses chosen to make the intermediate focus as large as possible to limit the intensity at the pinhole. However, implementing this at the required powers on the order of 100 W is challenging and would require careful optical and thermal design. Following this, the full chain from the red MOT through the dipole guide into the reservoir trap should be characterized and optimized in terms of transported atom flux, the stability of that flux, and temperature.

Towards continuous superradiance. With atoms recaptured in the reservoir, the next step is to load them into the moving optical lattice, formed by the bowtie cavity and transport them through the cavity mode. Superradiant emission then requires implementing an incoherent pumping scheme. The goal will first be to demonstrate normal-mode splitting and afterwards collective emission on the broader ${}^1S_0 \rightarrow {}^3P_1$ kHz transition, which is less demanding on flux and temperature and can give a higher superradiant output. Once collective emission is shown here, continuous superradiance on the mHz clock transition can be attempted. Achieving this would realize an active optical frequency reference whose stability is written into the atoms themselves, less sensitive to the cavity noise and free of the dead time that limit passive optical clocks.

BIBLIOGRAPHY

[1] A. Franzen. *ComponentLibrary: A Free Vector Graphics Library for Optics*. URL: https://www.gwoptics.org/ComponentLibrary/ (accessed: 2026-07-09).

[2] Bureau International des Poids et Mesures. *The International System of Units (SI Brochure)*. Tech. rep. BIPM, 2019.

[3] A. D. Ludlow et al. "Optical atomic clocks". In: *Reviews of Modern Physics* 87.2 (2015), pp. 637–701. DOI: 10.1103/RevModPhys.87.637.

[4] N. Dimarcq et al. "Roadmap towards the redefinition of the second". In: *Metrologia* 61.1 (2024), p. 012001. DOI: 10.1088/1681-7575/ad17d2.

[5] C. W. Chou et al. "Frequency Comparison of Two High-Accuracy Al^+ Optical Clocks". In: *Phys. Rev. Lett.* 104 (7 2010), p. 070802. DOI: 10.1103/PhysRevLett.104.070802.

[6] B. J. Bloom et al. "An optical lattice clock with accuracy and stability at the 10^{-18} level". In: *Nature* 506 (2014), pp. 71–75. DOI: 10.1038/nature12941.

[7] K. Numata, A. Kemery, and J. Camp. "Thermal-Noise Limit in the Frequency Stabilization of Lasers with Rigid Cavities". In: *Phys. Rev. Lett.* 93 (25 2004), p. 250602. DOI: 10.1103/PhysRevLett.93.250602.

[8] M. Gross and S. Haroche. "Superradiance: An essay on the theory of collective spontaneous emission". In: *Physics Reports* 93.5 (1982), pp. 301–396. DOI: 10.1016/0370-1573(82)90102-8.

[9] R. H. Dicke. "Coherence in Spontaneous Radiation Processes". In: *Phys. Rev.* 93 (1 1954), pp. 99–110. DOI: 10.1103/PhysRev.93.99.

[10] J. G. Bohnet et al. "A steady-state superradiant laser with less than one intracavity photon". In: *Nature* 484.7392 (2012), pp. 78–81. DOI: 10.1038/nature10920.

[11] H. M. Goldenberg, D. Kleppner, and N. F. Ramsey. "Atomic Hydrogen Maser". In: *Phys. Rev. Lett.* 5 (8 1960), pp. 361–362. DOI: 10.1103/PhysRevLett.5.361.

[12] D. Meiser et al. "Prospects for a millihertz-linewidth laser". In: *Phys. Rev. Lett.* 102 (2009), p. 163601. DOI: 10.1103/PhysRevLett.102.163601. arXiv: 0901.3105.

[13] S. L. Kristensen et al. "Subnatural Linewidth Superradiant Lasing with Cold ^{88}Sr Atoms". In: *Phys. Rev. Lett.* 130.22 (2023), p. 223402. DOI: 10.1103/PhysRevLett.130.223402.

[14] G. J. Dick. "Local Oscillator Induced Instabilities in Trapped Ion Frequency Standards". In: *Proceedings of the 19th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*. Redondo Beach, CA, 1987.

[15] H. Katori. "Longitudinal Ramsey spectroscopy of atoms for continuous operation of optical clocks". In: *Applied Physics Express* 14 (2021), p. 072006. DOI: 10.35848/1882-0786/ac0e16.

[16] M. A. Norcia et al. "Superradiance on the millihertz linewidth strontium clock transition". In: *Science Advances* 2.10 (2016), e1601231. DOI: 10.1126/sciadv.1601231.

[17] M. A. Norcia and J. K. Thompson. "Cold-Strontium Laser in the Superradiant Crossover Regime". In: *Phys. Rev. X* 6 (2016), p. 011025. DOI: 10.1103/PhysRevX.6.011025.

[18] S. B. Jäger et al. “Superradiant emission of a thermal atomic beam into an optical cavity”. In: *Phys. Rev. A* 104.3 (2021), p. 033711. DOI: 10.1103/physreva.104.033711.

[19] F. Famà et al. “Continuous cavity QED with an atomic beam”. In: *Phys. Rev. A* 110.6 (2024), p. 063721. DOI: 10.1103/physreva.110.063721.

[20] S. Bennetts et al. “Steady-State Magneto-Optical Trap with 100-Fold Improved Phase-Space Density”. In: *Phys. Rev. Lett.* 119 (22 2017), p. 223202. DOI: 10.1103/PhysRevLett.119.223202.

[21] C.-C. Chen et al. “Continuous Bose-Einstein condensation”. In: *Nature* 606.7915 (2022), pp. 683–687. DOI: 10.1038/s41586-022-04731-z.

[22] C. J. Foot. *Atomic Physics*. Vol. 7. Oxford Master Series in Physics. Oxford University Press, 2005. ISBN: 978-0-19-850695-9.

[23] M. Zhu, C. W. Oates, and J. L. Hall. “Continuous high-flux monovelocity atomic beam based on a broadband laser-cooling technique”. In: *Phys. Rev. Lett.* 67 (1 1991), pp. 46–49. DOI: 10.1103/PhysRevLett.67.46.

[24] E. L. Raab et al. “Trapping of Neutral Sodium Atoms with Radiation Pressure”. In: *Phys. Rev. Lett.* 59 (23 1987), pp. 2631–2634. DOI: 10.1103/PhysRevLett.59.2631.

[25] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. “Optical dipole traps for neutral atoms”. In: *Advances in Atomic, Molecular, and Optical Physics* 42 (2000), pp. 95–170. DOI: 10.1016/S1049-250X(08)60186-X. arXiv: physics/9902072.

[26] B. E. A. Saleh and M. C. Teich. “Beam Optics”. In: *Fundamentals of Photonics*. John Wiley & Sons, Ltd, 1991. Chap. 3, pp. 80–107. ISBN: 9780471213741. DOI: 10.1002/0471213748.ch3.

[27] T. A. Savard, K. M. O’Hara, and J. E. Thomas. “Laser-noise-induced heating in far-off resonance optical traps”. In: *Phys. Rev. A* 56 (2 1997), R1095(R)–R1098(R). DOI: 10.1103/PhysRevA.56.R1095.

[28] A. Kramida et al. *NIST Atomic Spectra Database (version 5.12)*. National Institute of Standards and Technology, Gaithersburg, MD. 2024. URL: https://physics.nist.gov/asd (accessed: 2026-05-29).

[29] S. B. Nagel et al. “Photoassociative Spectroscopy at Long Range in Ultracold Strontium”. In: *Phys. Rev. Lett.* 94 (8 2005), p. 083004. DOI: 10.1103/PhysRevLett.94.083004.

[30] M. M. Boyd. “High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard of Frequency and Time”. PhD thesis. Boulder, Colorado, USA: University of Colorado, 2007.

[31] P. Morzyński et al. “Absolute Measurement of the $^1S_0$–$^3P_0$ Clock Transition in Neutral ^{88}Sr over the 330 km-Long Stabilized Fibre Optic Link”. In: *Scientific Reports* 5.1 (2015), p. 17495. DOI: 10.1038/srep17495.

[32] S. Dubey et al. “Modeling of a continuous superradiant laser on the sub-mHz $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_0$ transition in neutral strontium-88”. In: *Phys. Rev. Res.* 7 (1 2025), p. 013292. DOI: 10.1103/PhysRevResearch.7.013292.

[33] S. Boughdachi et al. “Quasicontinuous sub-μK strontium source without a high-finesse cavity-stabilized laser”. In: *Phys. Rev. Applied* 25.3 (2026), p. 034053. DOI: 10.1103/3gnx-4s97.

[34] E. Peik et al. “Bloch oscillations and an accelerator for cold atoms”. In: *Applied Physics B* 65.6 (1997), pp. 685–692. DOI: 10.1007/s003400050334.

[35] A. V. Taichenachev et al. “Magnetic Field-Induced Spectroscopy of Forbidden Optical Transitions with Application to Lattice-Based Optical Atomic Clocks”. In: *Phys. Rev. Lett.* 96 (8 2006), p. 083001. DOI: 10.1103/PhysRevLett.96.083001.

[36] National Institute of Standards and Technology. *Atomic Weights and Isotopic Compositions with Relative Atomic Masses*. URL: https://www.nist.gov/pml/atomic-weights-and-isotopic-compositions-relative-atomic-masses (accessed: 2026-06-08).

[37] T. M. Brzozowski et al. “Time-of-flight measurement of the temperature of cold atoms for short trap-probe beam distances”. In: *Journal of Optics B: Quantum and Semiclassical Optics* 4 (2002), pp. 62–66. DOI: 10.1088/1464-4266/4/1/310.

[38] A. Kiruga et al. “Portal for high-precision atomic data and computation”. In: *Computer Physics Communications* 319 (2026), p. 109951. DOI: 10.1016/j.cpc.2025.109951.

[39] P. J. Mohr et al. “CODATA recommended values of the fundamental physical constants: 2022”. In: *Rev. Mod. Phys.* 97 (2 2025), p. 025002. DOI: 10.1103/RevModPhys.97.025002.

[40] P. A. Bélanger. “Beam propagation and the ABCD ray matrices”. In: *Opt. Lett.* 16.4 (1991), pp. 196–198. DOI: 10.1364/OL.16.000196.

[41] Q. Sheng et al. “Enhancing the field of view of cat-eye retroreflectors by simply matching the mirror radius of curvature and the lens focal length”. In: *Results in Physics* 37 (2022), p. 105558. DOI: 10.1016/j.rinp.2022.105558.

[42] B. Heizenreder et al. “Continuous cloud position spectroscopy using a magneto-optical trap”. In: *Phys. Rev. Applied* 25.3 (2026), p. L031003. DOI: 10.1103/wdp6-98wh.

[43] J. Bechhoefer. “Feedback for physicists: A tutorial essay on control”. In: *Rev. Mod. Phys.* 77 (3 2005), pp. 783–836. DOI: 10.1103/RevModPhys.77.783.

[44] J. G. Ziegler and N. B. Nichols. “Optimum settings for automatic controllers”. In: *Transactions of the American Society of Mechanical Engineers* 64 (1942), pp. 759–768.

[45] Thorlabs, Inc. *Photodiode Tutorial*. 2026. URL: https://www.thorlabs.com/photodiodes-and-photoconductors-tutorials (accessed: 2026-06-26).

[46] Thorlabs, Inc. *Optical Tables Tutorial: Vibrations*. URL: https://www.thorlabs.com/optical-tables-tutorial?tabName=Vibrations (accessed: 2026-07-09).

[47] R. Paschotta. *Shot Noise*. RP Photonics Encyclopedia. 2025. URL: https://www.rp-photonics.com/shot_noise.html (accessed: 2026-07-09).

ACKNOWLEDGMENTS

I gratefully acknowledge the Studienstiftung des deutschen Volkes for their financial support, which made it possible for me to do my Master's in Amsterdam.

First of all, I would like to thank my supervisor, Florian Schreck, for giving me the opportunity to work on this experiment. Your enthusiasm for the physics, superradiance in particular, is clear in the way you explain it, and I am grateful that you trusted us to present the experiment at several conferences.

Benedikt, thank you for spending so much time helping us on the experiment, even when you were busy with much more important things, and for always being happy to explain things in detail. Your patience and your optimism when new problems arise are things I'll try to learn. Ananya, thank you for the welcoming atmosphere you bring to the lab, for not once making me feel bad about not knowing something, and for including us Master students in things outside the lab. Both of you also showed me that you can somehow still make it fun, even when something is annoying and doesn't look very promising. I hope your hard work on the kHz experiment gets all the recognition it deserves.

Luuk, it was great having you around for the whole project, figuring out all the Master student things together. I really admire the focus you have, working for many hours on the same experimental problem without ever needing a break. Thanks to the rest of the group for always being happy to help: Alex, who went straight into solving any problem the moment it was mentioned; Ivo for the life advice; Digvijay for his lab equipment and his humor; Davide for the cappuccino breaks; and Yu Chih for the good company. Thanks also to András, who enthusiastically demonstrated and explained his PID project to me a year before I knew I would have my own, and to Ying and Haiwei, who make me excited for the experiment's future.

Thanks to the Technology Center of the UvA, in particular Edwin, for building the PID that worked so well and for helping me use and adapt the devices even when I didn't understand much about electronics.

Thanks to the IOP and everyone in it, for all the events and the truly nice atmosphere, with so many people trying to make it a friendly place to do physics. I'll miss that, and everyone I always run into in the hallways, especially from the ions lab and many of the soft matter people.

I am so grateful to have my AMEP friends, who always cheered me up when things didn't work, sometimes with unhinged conversations in the IOP lounge. Tommi, I could always come to you when I was excited, annoyed, bored, or hungry in the lab. You're a good listener, and you stay calm with complicated things in a way I wish I could.

And Valentin, my biggest fan, who had to deal with my frustration, my monologues about atomic physics, and convinced me many times that I could do this. Thank you!