



UNIVERSITÀ
DEGLI STUDI
FIRENZE

PhD in Physics and Astronomy
CYCLE XXXVIII

Coordinator Prof. Giovanni Modugno

CONTROL AND TRANSPORT OF SINGLE
ATOMS TRAPPED IN OPTICAL TWEEZERS
FOR QUANTUM SIMULATORS

CANDIDATE: ANDREA FANTINI

Supervisor : Prof. Leonardo Fallani

Co-supervisor : Dr. Vladislav Gavryusev

Anno Accademico 2025

CONTENTS

1	Introduction on Quantum Simulators	5
1.1	Historical Developments and State of the Art	7
1.1.1	Neutral Atoms in Optical Tweezers	9
1.2	Scope and Structure of this Thesis	10
2	Theoretical Description of Ultracold Atoms in Optical Tweezers	13
2.1	Cooling and Trapping Atoms: Atom–Light Interaction . .	13
2.1.1	Radiation pressure and optical forces	14
2.1.2	Doppler cooling and sub-doppler mechanisms . . .	15
2.1.3	Magneto-Optical Traps (MOTs)	21
2.1.4	Far off-resonant atom-light interaction	24
2.1.5	Tight trapping via Optical Tweezer	28
2.1.6	Pairwise losses via Light-Assisted Collisions	31
2.1.7	Ground state cooling via Sisyphus technique	33
2.2	Choosing the element: strontium 88	35
2.2.1	Electronic structure and level diagram	35
2.2.2	Key Transitions of Strontium	37
2.3	Tuning the interaction: Rydberg atoms	39
2.3.1	Rydberg physics: scaling laws and blockade	39
2.3.2	Optical excitation pathways	43
3	Architecture of the Experimental Machine	45
3.1	Hardware: the atom-light machine	46
3.1.1	Vacuum system and atomic source	46
3.1.2	Magnetic field control	50
3.1.3	Laser system and stabilization	52
3.1.4	Optical setup for 3D MOT trapping	58
3.2	Quantum platform: Optical Tweezer Generation	60
3.2.1	High-NA Objective	60
3.2.2	Acousto-Optic Deflectors (AODs)	61
3.2.3	Spatial Light Modulators (SLMs)	64
3.2.4	Tweezer generation setup	71
3.2.5	Atom Rearrangement and Defect-Free Arrays . . .	73
3.3	Software: Control System Development	75
3.3.1	Labscript Suite	75
3.3.2	Devices Implementation	83
3.3.3	FPGA-Based Control Hardware	84

3.3.4	MOGLabs QRF Driver	88
3.3.5	Basler ACE Pylon Camera	90
3.3.6	Hamamatsu ORCA Quest Camera	93
3.3.7	Spectrum AWG Card	97
3.3.8	Feedback Strategy	105
3.3.9	Independent Lab Processes	108
4	Characterization of Machine Performance	111
4.1	Two-stage MOTs of Strontium	111
4.1.1	First Stage of Trapping: Blue MOT	114
4.1.2	Bridging Stages: Broad-Band Red MOT	119
4.1.3	Second Stage of Trapping: Single-Frequency Red MOT	121
4.2	Single atoms in Optical Tweezers	124
4.2.1	Tweezer generation	124
4.2.2	Imaging scheme	127
4.2.3	Photon collection	128
4.2.4	From multiple to single atoms	130
4.2.5	Signal analysis	136
5	Case study: kinetically-induced frustration dynamics	147
5.1	Magnon-Hole Dynamics	148
5.2	Experimental Preparation, Evolution and Detection	157
5.2.1	Site-Resolved State Preparation	158
5.2.2	Quasi-Adiabatic Evolution to Low-Energy States	159
5.2.3	State Detection and Measurement Bases	160
5.3	Transverse Spin Correlation Analysis	163
5.4	Summary and Outlook	167
6	Conclusions and Future Perspectives	169
	Acknowledgements	173
	Bibliography	175

*"There is nothing like looking, if you want to find something.
You certainly usually find something, if you look,
but it is not always quite the something you were after."
— J.R.R. Tolkien*

INTRODUCTION ON QUANTUM SIMULATORS

Arrays of single neutral atoms trapped in optical tweezers have emerged as a versatile platform for quantum science, driving advances in quantum simulation, quantum computation, and precision metrology [1, 2, 3]. These systems feature a unique combination of precise atomic control and scalability, making them ideally suited to investigate complex many-body quantum phenomena. In recent years, significant progress has been achieved, including demonstrations of quantum simulation with Rydberg interactions [4, 5] and entangling operations between individually trapped atoms [6, 7].

Before getting specific, let me summarize here below when and why this research journey has begun and where we are now.

Understanding the behavior and dynamics of complex quantum systems is one of the most formidable challenges of modern physics. The difficulty arises from the exponential growth of computational complexity, which renders classical computing methods ineffective for most quantum many-body problems: for a system of n two-level quantum particles, the Hilbert space - *the abstract vector space that contains all possible quantum states of the system* - grows as 2^n , rendering exact numerical simulation impossible even for moderately sized systems. In fact, simulating the full quantum dynamics of only a few dozen particles already exceeds the memory capacity of current classical computers. [8]. Moreover, fundamental quantum phenomena such as entanglement and interference give rise to non-classical correlations that cannot be efficiently represented using classical bits and probabilistic models. Many problems in quantum physics fall into computational complexity classes that are inaccessible to classical algorithms, even in principle, such as the simulation of generic quantum many-body dynamics, ground-state energy estimation of strongly correlated systems, and sampling problems associated with

quantum circuits [9, 10, 11, 12]. These limitations fundamentally constrain our understanding of many-body phenomena such as high-temperature superconductivity, quantum magnetism, and non-equilibrium dynamics in strongly correlated systems.

This challenge was famously articulated by Richard Feynman in 1981:

“Nature isn’t classical, dammit, and if you want to make a simulation of nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem, because it doesn’t look so easy.”

Feynman proposed the idea of building a controllable quantum system specifically designed to emulate another quantum system - what we now call a *quantum simulator* [13]. He recognized that the same exponential complexity that hinders classical computation could be exploited advantageously: a quantum device naturally encodes exponentially large Hilbert spaces without requiring exponentially more physical resources. This insight laid the foundation for the concept of the quantum simulator.

A quantum simulator is a controllable quantum system engineered to reproduce the dynamics of other, often more complex or inaccessible, quantum systems [14]. By tailoring its parameters to emulate a target Hamiltonian, the simulator enables direct experimental access to many-body quantum physics. Such analog quantum simulators offer a powerful alternative to fully digital quantum computers, particularly in the near term, where scalable fault-tolerant architectures remain out of reach [15].

Let us consider a quantum system whose state is denoted by $|\phi\rangle$. Its time evolution from the initial state $|\phi(0)\rangle$ to $|\phi(t)\rangle$ is governed by the unitary operator

$$U = e^{-iHt/\hbar}, \quad (1.1)$$

where H is the system Hamiltonian. To investigate this system experimentally, one can design a controllable simulator that evolves according to

$$U' = e^{-iH't/\hbar}, \quad (1.2)$$

where H' is the Hamiltonian of the simulator. The procedure involves preparing an initial state $|\psi(0)\rangle$, applying the evolution U' , and measuring the final state $|\psi(t)\rangle$. If there exists a mapping between the two systems - that is, between $|\phi(0)\rangle \leftrightarrow |\psi(0)\rangle$, $|\phi(t)\rangle \leftrightarrow |\psi(t)\rangle$, and $U \leftrightarrow U'$ - then the simulator faithfully reproduces the dynamics of the original system. This principle forms the basis of quantum simulation and is illustrated schematically in Fig. 1.1.




Figure 1.1: Schematic representation of a quantum system and its corresponding quantum simulator. The quantum state $|\phi(0)\rangle$ evolves in time into $|\phi(t)\rangle$ via the unitary transformation $U = e^{-iHt/\hbar}$. The quantum simulator evolves from the state $|\psi(0)\rangle$ into $|\psi(t)\rangle$ via $U' = e^{-iH't/\hbar}$. The simulator is designed to reproduce the behavior of the simulated system through specific mappings: $|\phi(0)\rangle \leftrightarrow |\psi(0)\rangle$, $|\phi(t)\rangle \leftrightarrow |\psi(t)\rangle$, and $U \leftrightarrow U'$. While the simulated system may be uncontrollable or experimentally inaccessible, the simulator is fully controllable. The big arrows denote controllable operations; solid arrows describe the time evolutions, while dashed arrows indicate the correspondence between the states.

1.1 HISTORICAL DEVELOPMENTS AND STATE OF THE ART

The concept of quantum simulation, originally proposed by Richard Feynman, was later formalized by Seth Lloyd in 1996, who demonstrated that quantum systems can efficiently simulate the dynamics of any local quantum Hamiltonian [16].

Since then, experimental efforts toward realizing quantum simulators have advanced rapidly, propelled by both technological innovation and a deeper understanding of how to engineer controllable quantum systems. Over the past few decades, several distinct physical platforms have emerged - each characterized by its own advantages and challenges - and together they have shaped the development of modern quantum simulation.

Among the earliest platforms to enable many-body quantum simulations in the laboratory were systems of ultracold atoms confined in optical lattices. Enabled by breakthroughs in laser cooling and optical trapping, these systems offered a natural realization of lattice Hamiltonians such as the Bose–Hubbard and Fermi–Hubbard models. A milestone achievement was the observation of the superfluid–Mott insulator transition in bosonic gases, which demonstrated the ability of quantum simulators to access strongly correlated quantum phases [17].

Subsequent optical lattice experiments have provided unprecedented insight into quantum magnetism [18, 19], topological states [20], and out-of-equilibrium many-body dynamics [21, 22].

Trapped ions represent another highly successful platform. Enabled by progress in radio-frequency trapping, precise laser control, and advanced vacuum and cryogenic techniques, trapped-ion systems offer exceptional coherence times, reaching $T_2 \sim 1 - 100$ s for hyperfine qubits and on the order of 1 s for optical qubits, and allow high-fidelity initialization, manipulation, and readout of individual qubits [23, 24]. Through the collective vibrational modes of ion chains, effective spin models - such as transverse-field Ising Hamiltonians - can be realized with tunable interaction ranges [25, 26, 27].

Landmark experiments in this field have demonstrated quantum phase transitions, non-equilibrium dynamics, and the generation of entangled states involving tens of individually addressable ions [28, 29]. These achievements establish trapped ions as a versatile platform for both analog and digital quantum simulation schemes.

Superconducting circuits, based on Josephson junctions coupled to microwave resonators, have also played a central role in the evolution of quantum simulators. Advances in microfabrication and cryogenic engineering have enabled the realization of strongly interacting arrays of artificial atoms (qubits) with tunable parameters and fast gate operations [30].

These systems have opened new avenues for exploring quantum optics in circuit QED architectures [31, 32], and for simulating models such as lattice gauge theories [33] and spin Hamiltonians [34]. Other platforms, including photonic systems [35, 36] and semiconductor quantum dots [37, 38], are also being investigated for specific applications in quantum simulation. While these approaches offer unique advantages, challenges related to scalability and controllable interactions remain significant.

1.1.1 *Neutral Atoms in Optical Tweezers*

Within this landscape, arrays of neutral atoms trapped in optical tweezers have recently emerged as a promising, powerful and highly adaptable platform for quantum simulation. By using tightly focused laser beams to confine individual atoms, these systems offer an exceptional degree of spatial control, individual addressability, and dynamic reconfigurability - features that were previously unattainable in large-scale (~ 1000 atoms) quantum architectures [39].

Compared to traditional optical lattices, optical tweezer arrays provide far greater flexibility. Atoms can be arranged in arbitrary geometries, repositioned dynamically, and initialized with near-unity filling using active rearrangement protocols [40]. When excited to high-lying Rydberg states, they exhibit strong, tunable, and long-range interactions, enabling the programmable realization of spin models with customizable geometry, range, and connectivity [2].




Figure 1.2: Schematic representation of an array of optical tweezers (red beams) trapping individual atoms (blue spheres), realized by focusing laser beams by a high Numerical-Aperture objective (underneath).

These capabilities have been used to explore a wide variety of many-body phenomena, including quantum phase transitions [1], topological models [41], and quantum optimization problems [42].

A particularly promising avenue involves alkaline-earth (or alkaline-earth-like) atoms such as strontium (Sr) and ytterbium (Yb) [43, 44]. These species offer long-lived metastable states, narrow optical transitions, and rich internal structures, making them ideal candidates for high-precision

control and quantum information encoding. Their combined electronic and nuclear spin degrees of freedom enable novel schemes for quantum memory, manipulation, and hybrid applications involving optical clocks and precision metrology.

Recent experimental milestones include the creation of defect-free arrays containing thousands of individually trapped atoms, along with coherent control over their quantum evolution [45, 46]. Together with an expanding experimental toolkit - including local addressing, Rydberg excitation, and real-time geometric reconfiguration - optical tweezer arrays now stand at the forefront of programmable quantum simulation. They offer an unparalleled combination of control, scalability, and versatility for probing complex many-body quantum physics in both equilibrium and non-equilibrium regimes.

1.2 SCOPE AND STRUCTURE OF THIS THESIS

This thesis presents the research work carried out during the three years of my PhD, devoted to the construction and development of the first Italian platform for quantum simulation based on strontium Rydberg atoms trapped in optical tweezers, in the Department of Physics and Astronomy of the *University of Florence* under the lead of Prof. Leonardo Fallani. The project stands at the crossroads of atomic physics, quantum optics, and quantum information science, with the overarching goal of establishing a scalable and coherent quantum simulator capable of probing complex many-body quantum phenomena.

A substantial part of this work was focused on the design and implementation of the experimental control system that orchestrates the operation of the apparatus. In particular, I developed and deployed the first implementation of the `labscript` suite in our laboratory, integrating it into a neutral-atom experiment in Florence for the first time. This required creating a flexible and modular software infrastructure capable of coordinating lasers, imaging systems, timing electronics, AOD/SLM control, and data acquisition. Beyond the initial integration, I contributed major improvements to the system, enhancing its execution speed, stability, and accessibility for new users. These contributions significantly reduced experimental overhead, streamlined daily operation, and enabled a more intuitive and maintainable control environment for the entire group.

Beyond the technical implementation, this thesis also explores theoretical and computational aspects relevant to Rydberg-mediated quantum simulation. The results presented here thus reflect both the experimental realization of a new platform and the methodological groundwork required for its future applications. A case study is the work conducted during a four-month research stay at the *Institute d'Optique* in Paris, under the supervision of Prof. Antoine Browaeys, where I contributed to one of the latest experimental applications of Rydberg-atom simulators.

The structure of the thesis is as follows:

- Chapter 1 introduces the concept of quantum simulation and situates the work within the broader context of neutral-atom platforms.
- Chapter 2 discusses the theoretical framework underlying optical trapping and Rydberg interactions in strontium atoms.
- Chapter 3 describes the experimental setup, with emphasis on the optical architecture and the control system, and details the software infrastructure developed for experiment automation and device communication.
- Chapter 4 presents characterization measurements of the machine's performance and the first experimental results obtained with the platform.
- Chapter 5 shows the work carried out during my stay in Paris, illustrating one of the most recent applications of such platforms.
- Finally, Chapter 6 draws conclusions and outlines possible directions for future work.

THEORETICAL DESCRIPTION OF ULTRACOLD ATOMS IN OPTICAL TWEEZERS

The interaction between atoms and light is a cornerstone of modern atomic physics, especially in the fields of laser cooling and trapping. This chapter provides an overview of the key atom–light interaction concepts that underpin the experimental work presented in this thesis.

The material collected in this chapter serves as a concise and practical compendium of the theoretical foundations and technological tools typically covered in an ultracold-atoms course. It is intended to supply the essential background required to understand the physical principles behind our experimental approach, our objectives, and our results.

2.1 COOLING AND TRAPPING ATOMS: ATOM–LIGHT INTERACTION

To describe the interaction between atoms and light in an intuitive manner, it is common to employ the semiclassical approximation. In this framework, the internal degrees of freedom of the atom are treated quantum mechanically, while the electromagnetic field is modeled as a classical monochromatic wave. This description is accurate in the regime of large photon numbers, where quantum fluctuations of the radiation field can be neglected.

A simple yet highly instructive model within this approximation is the two-level atom driven by a near-resonant laser field. Within the rotating-wave approximation (RWA), the Hamiltonian can be written as:

$$H = \frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}\sigma_x, \quad (2.1)$$

where $\Delta = \omega_L - \omega_0$ is the detuning between the laser frequency ω_L and the atomic transition frequency ω_0 , Ω is the Rabi frequency (proportional to the electric-field amplitude), and σ_i are the Pauli matrices acting on

the two-level basis. The time evolution of the atomic state is governed by the optical Bloch equations, or equivalently by the master equation for the density matrix ρ :

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[H, \rho] + \mathcal{L}[\rho], \quad (2.2)$$

where the superoperator $\mathcal{L}[\rho]$ accounts for dissipative processes such as spontaneous emission [47, 48]. This formalism captures both transient and steady-state dynamics, describing not only population evolution but also coherences between the atomic states.

During absorption and emission processes, the atom exchanges momentum with the light field. Each absorption or spontaneous-emission event imparts a recoil momentum of $\hbar k$ to the atom, where k is the wavevector of the light. The cumulative effect of these scattering events gives rise to radiative forces - including the radiation pressure force and the optical dipole force - which are fundamental to laser cooling and optical trapping.

These concepts form the foundation for understanding the laser cooling mechanisms, radiation-pressure forces, and light-assisted trapping techniques discussed in the following sections. A fully quantum treatment of atom–light interactions (particularly in low-intensity fields or in optical cavities in strong coupling regime) requires quantization of the light field, for example using the Jaynes–Cummings model or general master-equation approaches. Nevertheless, for most laser-cooling and trapping applications, the semiclassical picture provides an excellent approximation.

2.1.1 *Radiation pressure and optical forces*

When an atom absorbs a photon from a near-resonant laser beam, it receives a momentum kick of magnitude $p = \hbar k$ in the direction of the wavevector k , where $k = 2\pi/\lambda$ and λ is the wavelength of the light. After a characteristic time determined by the excited-state lifetime, the atom spontaneously emits a photon in a random direction and returns to the ground state. Over many absorption–emission cycles, this process results in an average momentum transfer of $\hbar k$ per scattering event.

This mechanism gives rise to the radiation pressure force, a fundamental dissipative process that underlies laser cooling. For a two-level atom illuminated by a monochromatic plane wave, the average photon

scattering rate - that is, the rate at which the atom absorbs and re-emits photons - is given by:

$$\Gamma_{\text{sc}} = \frac{\Gamma}{2} \frac{s}{1 + s + \left(\frac{2\Delta}{\Gamma}\right)^2}, \quad (2.3)$$

where Γ is the natural linewidth of the atomic transition, $\Delta = \omega_L - \omega_0$ is the detuning between the laser frequency ω_L and the atomic resonance frequency ω_0 , and $s = I/I_{\text{sat}}$ is the saturation parameter. The saturation intensity I_{sat} is defined as

$$I_{\text{sat}} = \frac{2\pi^2\hbar c\Gamma}{3\lambda^3}, \quad (2.4)$$

where $\hbar$ is the reduced Planck constant and c is the speed of light.¹ The corresponding radiation pressure force acting on the atom is therefore

$$\mathbf{F}_{\text{rad}} = \hbar \mathbf{k} \Gamma_{\text{sc}}. \quad (2.5)$$

The scattering rate reaches its maximum value $\Gamma_{\text{sc}} = \Gamma/2$ for resonant excitation ($\Delta = 0$) in the strong-driving regime ($s \gg 1$). This sets a fundamental limit on the achievable scattering rate and explains why atomic transitions with large natural linewidths Γ are particularly effective for manipulating and cooling atoms moving at high velocities.

2.1.2 Doppler cooling and sub-doppler mechanisms

The effective detuning experienced by an atom can be evaluated by considering the Doppler shift associated with its motion. For an atom moving with velocity $\mathbf{v}$ and interacting with a laser field of frequency ω_L and wavevector $\mathbf{k}$, the frequency of the light in the atom's rest frame is given by

$$\omega' = \omega_L - \mathbf{k} \cdot \mathbf{v}, \quad (2.6)$$

where the term $\mathbf{k} \cdot \mathbf{v}$ represents the Doppler shift due to the projection of the atomic velocity along the propagation direction of the laser beam.

¹ Note: Some references adopt slightly different conventions for the saturation intensity. Here I_{sat} is defined for a closed two-level transition driven by resonant light; different polarization or multilevel configurations may lead to different numerical prefactors.

Consequently, the effective detuning of the light from the atomic resonance becomes

$$\Delta = \omega' - \omega_0 = \omega_L - \mathbf{k} \cdot \mathbf{v} - \omega_0, \quad (2.7)$$

where ω_0 is the transition frequency of the atom at rest. This velocity-dependent detuning plays a key role in phenomena such as Doppler cooling and sub-Doppler laser cooling, where the direction and magnitude of the force experienced by the atom depend on its motion relative to the laser field. Substituting the Doppler-shifted detuning (Eq. (2.7)) into the scattering rate expression (Eq. (2.3)) yields the velocity-dependent, or Doppler-shifted, scattering rate. When the atom and the laser beam propagate in opposite directions ($\mathbf{k} \cdot \mathbf{v} < 0$), the Doppler shift increases the effective frequency seen by the atom. To satisfy the resonance condition ($\Delta = 0$) and achieve maximum scattering, the laser frequency must therefore be tuned slightly below the atomic resonance, a configuration known as *red detuning*. Conversely, if the atom moves in the same direction as the laser beam ($\mathbf{k} \cdot \mathbf{v} > 0$), the Doppler shift reduces the effective frequency in the atom's rest frame. In this case, the laser must be tuned above resonance (*blue detuning*) to compensate for the shift and restore the resonant condition.

A typical atomic source used in cold-atom experiments consists of a thermal atomic beam, produced by heating a solid or metallic sample in an oven. The atoms effuse from the hot reservoir through a narrow aperture, forming a directed beam characterized by a broad, non-thermal velocity distribution. To decelerate these fast atoms, one exploits the radiation pressure exerted by a counter-propagating, red-detuned laser beam, as illustrated schematically in Fig. 2.1. The radiation pressure force acting on an atom is given by Eq. (2.5). Owing to the Doppler effect, the detuning experienced by an atom moving with velocity v_0 along the propagation axis of the atomic beam becomes

$$\Delta = \omega_L - \omega_0 - kv_0, \quad (2.8)$$

where k is the laser wavevector, ω_L the laser frequency, and ω_0 the atomic transition frequency. The scattering force reaches its maximum magnitude when the laser is resonant with the atomic transition in the atom's rest frame, i.e. when $\Delta = 0$. This defines the *resonance condition*:

$$\omega_L - \omega_0 = -kv_0. \quad (2.9)$$

The Doppler effect therefore modifies the effective detuning experienced by a moving atom and plays a crucial role in determining the efficiency of laser cooling and slowing mechanisms.




Figure 2.1: Schematic representation of the Zeeman slower and transverse cooling mechanism. Atoms emitted from the thermal oven form an effusive atomic beam with divergence θ . A counter-propagating laser beam (Zeeman slower beam) decelerates the atoms along the longitudinal axis, while transverse cooling beams reduce the atomic divergence, thereby increasing the flux of slow atoms reaching the trapping region.

Another important phenomenon that influences the resonance condition of atomic transitions is the *Zeeman effect*, which arises from the interaction between the atomic magnetic dipole moment $\boldsymbol{\mu}$ and an external magnetic field $\mathbf{B}$. The corresponding interaction Hamiltonian can be expressed as

$$H_Z = -\boldsymbol{\mu} \cdot \mathbf{B}. \quad (2.10)$$

The total magnetic dipole moment $\boldsymbol{\mu}$ of an atom in a *LS coupling* regime can be expressed as

$$\boldsymbol{\mu} = -\frac{\mu_B}{\hbar} (g_L \hat{\mathbf{L}} + g_S \hat{\mathbf{S}}) - \frac{\mu_N}{\hbar} g_I \hat{\mathbf{I}}, \quad (2.11)$$

where $\hat{\mathbf{L}}$, $\hat{\mathbf{S}}$, and $\hat{\mathbf{I}}$ are the orbital, spin, and nuclear angular momentum operators, respectively. Here, μ_B and μ_N denote the Bohr and nuclear magnetons, while g_L , g_S , and g_I are the corresponding Landé g -factors.

For a magnetic field applied along the z axis, the Zeeman Hamiltonian reduces to the scalar form

$$\hat{H}_Z = \mu_B B (g_L \hat{L}_z + g_S \hat{S}_z) + \mu_N B g_I \hat{I}_z = \mu_B B g_J \hat{J}_z + \mu_N B g_I \hat{I}_z, \quad (2.12)$$

where $\hat{J}_z = \hat{L}_z + \hat{S}_z$ is the projection of the total electronic angular momentum operator $\hat{\mathbf{J}}$, and g_J is the Landé factor defined as

$$g_J = \frac{g_L [J(J+1) + L(L+1) - S(S+1)] + g_S [J(J+1) - L(L+1) + S(S+1)]}{2J(J+1)}. \quad (2.13)$$

For the bosonic isotopes of strontium, ^{88}Sr and ^{86}Sr , the nuclear spin is zero ($I = 0$); hence, the nuclear term in Eq. (2.12) vanishes, leaving only the electronic contribution. This is not the case for the fermionic isotope ^{87}Sr , where $I = 9/2$, or for other atoms possessing a non-zero hyperfine structure.

Finally, we note that for electronic states with total angular momentum $J = 0$, such as the ground state $^1S_0$ of strontium, the electronic Zeeman contribution vanishes identically since $\hat{J}_z = 0$. In this case, the Zeeman shift arises solely from the nuclear magnetic moment. However, the associated energy scale is strongly suppressed because the nuclear magneton μ_N is much smaller than the Bohr magneton μ_B ($\mu_N/\mu_B \simeq 1/1836$). As a consequence, even for fermionic isotopes with nonzero nuclear spin, the Zeeman sensitivity of $J = 0$ states remains extremely weak compared to states with finite electronic angular momentum.

In this regime, the Zeeman interaction leads to a shift of the atomic energy levels proportional to the magnetic field strength and the magnetic quantum number m_J , resulting in a linear splitting of the Zeeman sublevels:

$$\Delta E = \mu_B B g_J m_J. \quad (2.14)$$

Transverse Cooling

Atoms emerging from an atomic oven form a directed beam characterized by a broad velocity distribution and a finite angular divergence. The angular spread arises from the thermal motion of atoms inside the oven, which causes them to exit with varying transverse velocity components relative to the mean propagation direction. To reduce this divergence and improve beam collimation, *transverse laser cooling* can be applied.

In this technique, two counter-propagating laser beams are aligned along a direction orthogonal to the atomic beam propagation axis (for instance, along the x axis), as illustrated in Figure 2.1. Their radiation pressure exerts a velocity-dependent force that damps the transverse motion of the atoms.

Considering an atom moving at an angle θ with respect to the main propagation axis, the radiation pressure force experienced in the atomic reference frame, derived from Equations 2.3 and 2.5, can be expressed as

$$F_x = \frac{\hbar k \Gamma}{2} s \left[-\frac{1}{1 + s + \frac{4(\Delta + kv \sin \theta)^2}{\Gamma^2}} + \frac{1}{1 + s + \frac{4(\Delta - kv \sin \theta)^2}{\Gamma^2}} \right], \quad (2.15)$$

where k is the laser wavenumber, Γ is the natural linewidth of the transition, and Δ is the laser detuning in the laboratory frame. This expression accounts for the contributions from both cooling beams, propagating in opposite directions along x .

Transverse cooling affects the velocity components orthogonal to the atomic beam direction, thereby reducing the angular divergence of the atomic beam and increasing the flux of slow atoms that reach the trapping region. The relevant velocity component for transverse cooling is the projection of the atomic velocity along the cooling axis, given by $v \sin(\theta)$. When $\theta > 0$, the Doppler term $kv \sin(\theta)$ is positive. By selecting a red detuning ($\Delta < 0$), the term $(\Delta + kv \sin(\theta))^2$ in the first denominator of Eq. (2.15) becomes smaller than the corresponding term $(\Delta - kv \sin(\theta))^2$ in the second. This imbalance between the radiation pressures of the two counter-propagating beams generates a net damping force that pushes the atoms back toward $\theta = 0$, thereby reducing the angular divergence of the atomic beam.

In our experimental implementation, this same mechanism is also exploited to *deflect* the atomic beam emerging from the oven by an angle $\theta_d = 20^\circ$. Two nearly orthogonal laser beams, retroreflected by in-vacuum optics (see Chapter 3), create an intentionally unbalanced cooling configuration that introduces a controlled angular deviation of the mean atomic trajectory. The retroreflected beams are only partially reflected, resulting in different optical powers along the two propagation directions and hence in an asymmetric scattering force responsible for the desired deflection.

Zeeman Slower

Atoms coming out of a thermal oven must be decelerated before they can be captured in the subsequent cooling and trapping stages of the experiment. A widely used technique in ultracold atom physics combines

the *Doppler effect* and the *Zeeman effect* to continuously slow an atomic beam from initial velocities of several hundreds of meters per second down to a few tens of meters per second, effectively reducing the mean kinetic energy per atom.

Let us consider an atomic beam propagating along the positive y axis and a counter-propagating laser beam along the negative y direction (so that $-\mathbf{k} \cdot \mathbf{v} = kv$), in the presence of a magnetic field B_y oriented along the same axis. Since ^{88}Sr ground state has $J = 0$ (therefore $m_J = 0$), combining Equations 2.3 and 2.5, the resulting radiation pressure force along the beam axis can be written as

$$F_y = -\frac{\hbar k \Gamma}{2} \frac{s}{1 + s + \frac{4}{\Gamma^2} \left[\Delta + \frac{m'_J g'_J \mu_B B_y}{\hbar} + kv \right]^2}, \quad (2.16)$$

where g_J and g'_J denote Landé g -factors of the ground state and the excited state, while m'_J corresponds to the magnetic quantum number of the excited state. The difference between the Zeeman shifts of the two states, proportional to $m'_J g'_J \mu_B B_y$, determines the effective detuning of the transition in the presence of the magnetic field. The saturation parameter s is assumed constant and approximately uniform for a well-collimated laser beam.

As atoms are decelerated, the Doppler shift kv decreases; to maintain resonance and thus a constant scattering rate, a spatially varying magnetic field $B_y(y)$ is applied along the beam axis. A common design strategy is to impose a nearly constant deceleration a (as $a > 0$), which corresponds to maintaining a constant effective detuning in the atomic rest frame:

$$\Delta_{\text{atom-frame}} = \Delta + \frac{m'_J g'_J \mu_B B_y(y)}{\hbar} + kv(y) = \text{const.} \quad (2.17)$$

Maximizing the scattering rate corresponds to setting this detuning to zero, yielding the *resonance condition*:

$$-\Delta = \frac{m'_J g'_J \mu_B B_y(y)}{\hbar} + kv(y). \quad (2.18)$$

By defining

$$\xi_y = \frac{m'_J g'_J \mu_B}{\hbar}, \quad (2.19)$$

the spatial profile of the magnetic field that satisfies the resonance condition (Eq. 2.18) can be expressed as

$$B_y(y) = -\frac{\Delta}{\xi_y} - \frac{k}{\xi_y} \sqrt{v_0^2 + 2ay}, \quad (2.20)$$

where v_0 is the initial atomic velocity and a is the desired constant deceleration. The maximum achievable deceleration for a transition with linewidth Γ is given by

$$a_{\max} = \frac{\hbar k \Gamma}{2M}, \quad (2.21)$$

which scales with Γ , making broad atomic transitions preferable for efficient Zeeman slowing. For the blue transition of strontium ($\Gamma/2\pi = 32$ MHz), the maximum deceleration is $a_{\max} \simeq 10^6 \text{ m s}^{-2}$, requiring a Zeeman slower of length $L \simeq 10$ cm to decelerate atoms from $v_0 \simeq 450 \text{ m s}^{-1}$ to rest. A typical Zeeman slower is realized using a solenoid composed of concentric coils with a varying number of turns along the atomic beam axis. In our experiment, instead, a commercial atomic source (see Chapter 3) provides a set of permanent magnets approximately 5 cm long to generate the slowing field. This represents a practical, though non-optimal, solution, and the exact field profile is not disclosed due to patent restrictions.

The maximum initial velocity that can be slowed to a few tens of meters per second over the 5 cm interaction region is about $v_0 \simeq 320 \text{ m s}^{-1}$, even assuming $s \gg 1$. Furthermore, the assumption of a constant saturation parameter s is only an approximation, since the manufacturer recommends a slightly focused slowing laser beam to optimize performance. This results in a position-dependent radiation pressure force, which is stronger where the laser intensity is higher.

2.1.3 Magneto-Optical Traps (MOTs)

Atoms slowed by a Zeeman slower still retain residual kinetic energy and exhibit finite spatial and velocity distributions. To confine and further cool these atoms for precision experiments, it is necessary to perform a *Magneto-Optical Trap* (MOT). A MOT combines counter-propagating, red-detuned laser beams with a spatially varying magnetic field to generate both a velocity-dependent damping force (cooling) and a position-dependent restoring force (trapping). This configuration allows atoms to be simultaneously cooled and trapped near the trap center, forming a dense, ultracold sample suitable for subsequent experimental manipulation.

In most experiments, a 3D MOT configuration is used, where three orthogonal pairs of counter-propagating laser beams provide cooling along all spatial directions, while a quadrupole magnetic field centered at the beam intersection ensures spatial confinement. The atomic motion

along each Cartesian axis can be analyzed independently using scalar equations. The quadrupole field is typically generated using two coils in an anti-Helmholtz configuration (coaxial circular coils with equal currents flowing in opposite directions). The resulting magnetic field, in the intermediate point between the coils on their axis, can be expressed as

$$\mathbf{B} = b [-x, -y, 2z], \quad (2.22)$$

where b is the magnetic field gradient along z ($b > 0$). The factor of two in the z component and the opposite signs in the x and y components arise from Maxwell's condition $\nabla \cdot \mathbf{B} = 0$. Apart from these amplitude and sign differences, the three spatial directions are treated equivalently.




Figure 2.2: Zeeman splitting of the excited state $J' = 1$ in the presence of a magnetic field gradient. Two counter-propagating beams with opposite circular polarizations drive σ^- transitions from the non-degenerate ground state $J = 0$. Because the magnetic field reverses direction at $z = 0$, the quantization axis is defined locally along the magnetic field. Consequently, the identification of the $m_J' = \pm 1$ sublevels changes sign across the trap center, and the Zeeman branches are not continuous between the two sides of the diagram. On either side of the zero, the corresponding circularly polarized beam drives a σ^- transition with respect to the local magnetic-field direction.

We consider optical transitions from a ground state with total angular momentum $J = 0$ to an excited state with $J' = 1$, as is the case for the cooling transitions used in both the blue and red MOTs of ^{88}Sr . When the quadrupole magnetic field is applied, the energies of the excited-state Zeeman sublevels $|J' = 1, m_J' = \pm 1, 0\rangle$ shift linearly with the local magnetic field (Fig. 2.2).

The ground state $|J = 0\rangle$ remains unshifted since it has no magnetic moment. The two counter-propagating beams along the z -axis must be circularly polarized with opposite helicities. With the quantization axis defined locally along the magnetic-field direction, the corresponding circularly polarized beam drives a σ^- transitions ($\Delta m_J = -1$) on either side of the magnetic-field zero.

Because of the Zeeman effect, atoms displaced from the trap center experience a position-dependent shift in resonance frequency. For $z > 0$, where $B_z > 0$, the $m_J' = -1$ sublevel is shifted down in energy, bringing the σ^- beam (propagating toward $-z$) closer to resonance. Conversely, for $z < 0$, the opposite Zeeman shift favors excitation by the σ^+ -polarized beam. This behavior follows directly from the electric-dipole selection rules, which impose $\Delta m_J = \pm 1$ for $\sigma^\pm$ transitions, respectively.

This asymmetry in scattering rates between the two circular polarizations produces a restoring force toward the magnetic-field zero at the trap center. Combined with the Doppler-dependent frictional force from the red-detuned light, this mechanism yields both spatial confinement and velocity damping.

The radiation pressure force along the z -axis, including contributions from both counter-propagating beams, can be expressed as

$$F_z = \frac{\hbar k \Gamma}{2} \left[\frac{s}{1 + s + 4(\Delta - \delta_Z - kv)^2 / \Gamma^2} - \frac{s}{1 + s + 4(\Delta + \delta_Z + kv)^2 / \Gamma^2} \right], \quad (2.23)$$

where

$$\delta_Z = \frac{m_J' g_J' \mu_B b z}{\hbar}. \quad (2.24)$$

For small velocities ($kv \ll \Gamma$) and small displacements ($\delta_Z \ll \Gamma$), Eq. (2.23) can be linearized as

$$F_z(v, z) \simeq -\alpha v - \kappa z, \quad (2.25)$$

where the damping and restoring coefficients are

$$\alpha = 8\hbar k^2 \frac{(I/I_{\text{sat}})}{(1 + s + 4\Delta^2/\Gamma^2)^2} \frac{|\Delta|}{\Gamma}, \quad \kappa = \alpha \frac{g_J m_J \mu_B b}{\hbar k}. \quad (2.26)$$

This corresponds to a damped harmonic oscillator:

$$\frac{d^2 z}{dt^2} + \gamma \frac{dz}{dt} + \omega_{\text{trap}}^2 z = 0, \quad (2.27)$$

where $\gamma = M\alpha$ and $\omega_{\text{trap}} = \sqrt{\kappa/M}$, with M denoting the atomic mass. The MOT parameters can be tuned to optimize atom number, density, or capture velocity range depending on experimental priorities.

Because atoms continuously absorb and spontaneously re-emit photons, their motion exhibits a random walk in momentum space, setting a lower bound on the achievable kinetic energy and thus the final temperature. The Doppler cooling limit is given by

$$k_B T_D = \frac{\hbar \Gamma}{4} \left(1 + \frac{I}{I_{\text{sat}}} + \frac{4\Delta^2}{\Gamma^2} \right) \frac{2|\Delta|}{\Gamma}, \quad (2.28)$$

where Δ is the laser detuning and Γ the natural linewidth. The minimum Doppler temperature occurs for $\Delta = -\Gamma/2$:

$$T_D^{\text{min}} = \frac{\hbar \Gamma}{2k_B}. \quad (2.29)$$

The Doppler limit scales with the transition linewidth. For the broad $^1S_0 \leftrightarrow ^1P_1$ transition of strontium ($\Gamma/2\pi \simeq 32$ MHz), $T_D^{\text{min}} \simeq 720$ μK , whereas the narrow $^1S_0 \leftrightarrow ^3P_1$ transition ($\Gamma/2\pi \simeq 7.6$ kHz) yields ideally $T_D^{\text{min}} \simeq 180$ nK. While the broad transition efficiently captures atoms with velocities of a few m/s, the final temperature is too high for direct loading into optical tweezers. Hence, a two-stage MOT is often implemented: the first stage uses the broad transition for efficient capture, followed by a narrow-line MOT that cools the atoms to the μK regime and further-compacts them, making them suitable for precision optical trapping and quantum control experiments.

2.1.4 Far off-resonant atom-light interaction

We now introduce key theoretical aspects of atom-light interaction in the far off-resonant regime, which is crucial for understanding optical trapping with Optical Tweezers (O.T.). In this context, the interaction between an atom and an oscillating electromagnetic field gives rise to the *optical dipole potential*, also known as the *AC Stark shift* [49, 50, 51]. This interaction creates a conservative potential capable of trapping neutral atoms.

The optical dipole potential generated by a far off-resonant laser beam can be expressed in scalar approximation as [52]:

$$U(\mathbf{r}, z) = -\frac{1}{2\epsilon_0 c} I(\mathbf{r}, z) \alpha(\lambda_L), \quad (2.30)$$

where ϵ_0 is the vacuum permittivity, c is the speed of light, $I(\mathbf{r}, z)$ is the local intensity of the laser field at transverse position $\mathbf{r}$ and axial position

z , and $\alpha(\lambda_L)$ is the scalar atomic polarizability (generally a tensor) at the laser wavelength λ_L .

For a focused Gaussian beam, referred to here as an *optical tweezer*, the intensity profile is given by:

$$I(\mathbf{r}, z) = \frac{2}{\pi} \frac{P_L}{w^2(z)} \exp \left[-\frac{2r^2}{w^2(z)} \right], \quad (2.31)$$

where P_L is the laser power and $w(z)$ is the $1/e^2$ beam radius at position z . Substituting Eq. (2.31) into Eq. (2.30) gives:

$$U(\mathbf{r}, z) = -\frac{P_L \alpha(\lambda_L)}{\epsilon_0 c \pi w^2(z)} \exp \left[-\frac{2r^2}{w^2(z)} \right]. \quad (2.32)$$

The beam radius $w(z)$ varies along the propagation axis and reaches its minimum w_0 at the beam waist $z = 0$. Its axial dependence is described by

$$w(z) = w_0 \sqrt{1 + \frac{z^2}{z_R^2}}, \quad (2.33)$$

where $z_R = \pi w_0^2 / \lambda_L$ is the Rayleigh length, defining the characteristic scale over which the beam diverges.

If the light is red-detuned ($\omega_L < \omega_0$), the polarizability is positive, $\alpha(\lambda_L) > 0$, and Eq. (2.32) describes a potential well. Two key experimental parameters for optical dipole traps are the *trap depth* and the *trap frequencies*. The trap depth is defined as the maximum absolute value of the potential, occurring at the trap center ($r = 0, z = 0$):

$$U_0 = \frac{P_L \alpha(\lambda_L)}{\epsilon_0 c \pi w_0^2}. \quad (2.34)$$

In the regime where the trap depth is much greater than the atomic thermal energy, the potential near the center can be approximated as harmonic.

Expanding the Gaussian profile to second order for small radial displacements $r \ll w_0$ at $z = 0$ gives:

$$\exp \left[-\frac{2r^2}{w_0^2} \right] \approx 1 - \frac{2r^2}{w_0^2}. \quad (2.35)$$

Substituting this into Eq. (2.32) yields the *Harmonic Oscillator* (HO) approximation for the radial potential:

$$U_{\text{HO}}(r, 0) \approx -U_0 + \frac{2P_L \alpha(\lambda_L)}{\epsilon_0 c \pi w_0^4} r^2. \quad (2.36)$$

Equating the quadratic term to the harmonic oscillator potential $\frac{1}{2}M\omega_r^2 r^2$ gives the radial trap frequency:

$$\omega_r = \sqrt{\frac{4P_L \alpha(\lambda_L)}{M\epsilon_0 c \pi w_0^4}}. \quad (2.37)$$

Similarly, the axial potential near the trap center can be approximated as harmonic by expanding Eq. (2.32) along z at $r = 0$:

$$U(0, z) \approx -U_0 + \frac{P_L \alpha(\lambda_L)}{\epsilon_0 c \pi w_0^2 z_R^2} z^2, \quad (2.38)$$

which yields the axial trap frequency:

$$\omega_z = \sqrt{\frac{2P_L \alpha(\lambda_L)}{M\epsilon_0 c \pi w_0^2 z_R^2}}. \quad (2.39)$$

Both the trap depth and trap frequencies strongly depend on the optical tweezer parameters. The trap depth scales linearly with laser power and inversely with the square of the waist, while the radial trap frequency scales as the square root of the power and inversely with the square of the waist. In the axial direction, the Rayleigh length introduces an additional dependence, but since $z_R \propto w_0^2$, the overall axial frequency scales approximately inversely with the cube of the waist. These scaling properties enable micrometer-sized optical tweezers to generate deep traps with only a few milliwatts of laser power per site, a key feature for extending the system to large arrays.

Atomic polarizability

The optical dipole potential can act either as an attractive or repulsive force on an atom, depending on the internal structure of the atom and the sign of its polarizability. In general, the light shift induced by the interaction is largest where the laser intensity is highest. If this shift lowers the energy of the atomic state, the atom is attracted toward regions of maximum intensity and experiences a confining potential. Conversely, if the shift increases the energy, the atom is repelled from the light field, experiencing an anti-trapping effect.

Calculating the polarizability $\alpha(\lambda_L)$ for an off-resonant trapping wavelength λ_L can be non-trivial, particularly because the two-level approximation often fails in realistic atomic systems. When the laser is far detuned from multiple transitions, the total polarizability must account

for contributions from all relevant excited states. In such cases, a useful expression for the dynamic polarizability of an electronic state i is given by [53]:

$$\alpha_i(\omega_L) = \frac{6\pi\epsilon_0 c^3}{\omega_{ik}^3} \sum_k \frac{A_{ik}}{\omega_{ik}^2 - \omega_L^2}, \quad (2.40)$$

where $\omega_L = 2\pi c/\lambda_L$ is the angular frequency of the trapping laser, A_{ik} are the Einstein A coefficients (proportional to the square of the dipole matrix elements), and ω_{ik} are the transition frequencies between the initial state i and the excited states k . Polarizabilities generally depend on the polarization of the trapping light, as well as on the atomic fine and hyperfine structures, but in our treatment they have not been taken into account. External magnetic fields may also affect the polarizability by shifting magnetic sublevels with non-zero magnetic moments. For detailed computational methods and tabulated values, see Refs. [6, 54]. In particular, Ref. [55] provides a comprehensive approach for including fine and hyperfine interactions as well as Zeeman effects in polarizability calculations for Strontium atoms.

State	$ m_J $	Polarizability (a.u.)
$5s^2 \ ^1S_0$	0	286
$5s5p \ ^3P_0$	0	286
$5s5p \ ^3P_1$	0	195.1
	1	353.9
$5s5p \ ^3P_2$	0	464.3
	1	391.6
	2	173.8
$5s4d \ ^1D_2$	0	256.8
	1	522.2
	2	1318
$5s5p \ ^1P_1$	0	178.4
	1	737.8
$5s6s \ ^3S_1$	0	-494.1
	1	-481.9

Table 1: Summary of strontium polarizabilities at $\lambda_L = 813.4$ nm, assuming linear polarization along $\hat{x}$ and zero magnetic field. Adapted from Ref. [6].

A central concept in optical trapping is the *magic wavelength* condition. This occurs when two atomic levels have identical polarizabilities at a specific wavelength λ_{magic} , resulting in equal energy shifts for both states. At this wavelength, the differential light shift vanishes, and the transition frequency between the two levels remains unaffected by the trapping field, independent of intensity. Non-magic traps, in contrast, introduce light shifts that vary with laser power, potentially degrading coherence since power fluctuations and inhomogeneities can lead to infidelities when driving transitions. Magic trapping conditions are therefore essential in precision measurement platforms, especially when dealing with narrow optical transitions, as in optical atomic clocks.

Strontium, in particular, exhibits a magic wavelength at $\lambda_{\text{magic}} = 813.4$ nm for the clock transition $^1S_0 \rightarrow ^3P_0$. This wavelength lies in the infrared, where high-power Ti:Sapphire lasers are readily available. Calculated polarizability values for key Strontium states at $\lambda_L = 813.4$ nm are provided in Ref. [6] and summarized in Table 1.

2.1.5 Tight trapping via Optical Tweezer

An optical tweezer is created by tightly focusing a laser beam. This can be achieved with a simple lens, but far superior performance is obtained with a high-quality microscope objective. Since these traps operate on micrometer length scales, the optical system must be designed with great care to suppress aberrations that would otherwise distort the trapping potential. For this reason, diffraction-limited and aberration-corrected objectives are typically used. Compared to standard achromatic lenses, they provide higher performance and can be optimized simultaneously for multiple wavelengths. This capability is essential, as the same objective is used not only to generate the optical tweezer but also to collect fluorescence and form site-resolved images of the trapped atoms (see Sec. 4.2.5).

A laser beam can be treated as a Gaussian beam provided that the paraxial approximation holds, meaning that the amplitude A of the electromagnetic field, propagating along the z direction with wavelength λ , varies slowly compared to the optical wavelength [56]:

$$\frac{\partial A}{\partial z} \ll kA, \quad (2.41)$$

where $k = 2\pi/\lambda$ is the wave number. Under this approximation, the electric field

$$E(r, z, t) = A(r, z)e^{i(\omega t - kz)}$$

takes the form:

$$A(r, z) = A_0 \frac{w_0}{w(z)} \exp\left[-\frac{r^2}{w^2(z)}\right] \exp\left[-ikz - i\frac{kr^2}{2R(z)} + i\zeta(z)\right], \quad (2.42)$$

where A_0 is the peak amplitude, r is the radial coordinate, $\zeta(z) = \arctan(z/z_R)$ is the Gouy phase and w_0 is the waist radius at the focal plane as defined in Eq. 2.33 with $w(z)$. The radius of curvature of the wavefront is

$$R(z) = z \left[1 + \left(\frac{z_R}{z} \right)^2 \right]. \quad (2.43)$$

For a tightly focused beam, as in an optical tweezer, the paraxial approximation generally fails because the field varies significantly on length scales comparable to z_R , often only a few micrometers and thus similar to the wavelength. The limits of the paraxial approximation and the resulting higher-order corrections are discussed in detail in Ref. [57], including their connection to optical aberrations.




Figure 2.3: Airy disk intensity profile in two and one dimensions. The central bright peak is surrounded by concentric rings originating from diffraction due to the finite aperture of the objective.

Even in a perfectly aberration-free optical system, the image of a point-like emitter has a finite spatial extent due to diffraction. The response of the imaging system to a point source is described by the point spread function (PSF). For an ideal diffraction-limited system with a circular pupil of diameter $D = 2a$, the PSF is given by the Airy pattern (Fig. 2.3) [58].

In the far field, the intensity distribution as a function of the observation angle θ relative to the optical axis reads

$$I(\theta) = I_0 \left[\frac{2J_1(ka \sin \theta)}{ka \sin \theta} \right]^2, \quad (2.44)$$

where J_1 is the first-order Bessel function of the first kind, $k = 2\pi/\lambda$ is the optical wavenumber, and a is the radius of the circular aperture.

The radius of the first minimum of the Airy pattern defines the conventional spot size limit, obtained when

$$ka \sin \theta \approx 1.22 \pi. \quad (2.45)$$

For an optical system with numerical aperture NA,

$$\text{NA} = n \sin \theta_{\max}, \quad (2.46)$$

where $\theta_{\max} = \arctan(D/2f)$ is the maximum aperture angle (with focal length f), and n is the refractive index ($n = 1$ in vacuum). This leads to the diffraction-limited radius

$$a_0 = \frac{0.61\lambda}{\text{NA}}, \quad (2.47)$$

which is derived under the assumption of a uniformly illuminated pupil, corresponding to a plane wave incident on the back focal plane (BFP) of the objective and resulting in the image of a point source. In the presence of a Gaussian illumination at the BFP, as is typically the case for tightly focused laser beams, the effective PSF may differ slightly from the ideal Airy pattern, introducing a correction factor of order unity.

Real imaging systems typically exhibit a larger PSF than the theoretical diffraction limit due to residual optical aberrations. The best achievable resolution is therefore characterized by the measured Airy radius r_0 . If $r_0 = a_0$, the system is said to be diffraction limited. In practice, this condition is achieved only within a finite field of view (FOV) around the optical axis and typically only at specific design wavelengths.

Achieving the smallest possible PSF is crucial in our experiment, both to generate tightly confining optical tweezers—yielding large trap depths and trapping frequencies—and to ensure high-resolution imaging. The objective used in this work (see Sec. 3.2.1) is designed to operate at the diffraction limit within its specified FOV at both 813 nm (tweezers) and 461 nm (fluorescence imaging), with a numerical aperture NA= 0.55. The corresponding diffraction-limited spot radii are

$$a_0(461, \text{nm}) = 502 \text{ nm}, \quad a_0(813, \text{nm}) = 902 \text{ nm}. \quad (2.48)$$

2.1.6 Pairwise losses via Light-Assisted Collisions

A key feature of optical tweezer platforms is their ability to isolate a single atom in each trap. To load atoms into the optical dipole traps, we spatially overlap the tweezers with the cold atomic cloud produced by the narrow-line red MOT described in Sec. 2.1.3. The traps typically have depths on the order of tens to hundreds of μK , whereas the atoms in the MOT have a thermal energy of only about $5 \mu\text{K}$.

The number of atoms loaded from the MOT into a single tweezer follows a Poisson distribution [59]. However, extracting the mean atom number directly from fluorescence imaging is inherently challenging, as it is subject to systematic errors. This is due to two main limitations: first, the detection camera can reliably distinguish between zero atoms and one-or-more atoms, but it cannot easily resolve multiple occupancies²; second, the imaging process itself can induce atom loss, altering the atom number during the measurement.

For a Poisson distribution with mean atom number $\tilde{N}$, the probability of loading zero atoms is $P(0) = e^{-\tilde{N}}$. Experimentally, the fraction of images corresponding to empty tweezers provides an upper bound on $P(0)$, from which we obtain a lower bound on the mean occupation number, $\tilde{N} = -\ln(P(0))$.

One possible strategy to achieve primarily single-atom loading is to carefully tune the red MOT and trap parameters to maximize the probability of loading exactly one atom. However, Poissonian statistics fundamentally limit this approach: the maximum achievable probability of loading a single atom is only about 37%, with another 37% of tweezers remaining empty and roughly 26% containing multiple atoms [60].

To prepare exactly at most one atom per trap, a widely used mechanism in single-atom experiments is pairwise loss via light-assisted collisions (LACs) [61, 62, 63]. When two atoms occupy the same trap and are illuminated, near-resonant light can excite the atom pair into an attractive molecular potential. The atoms accelerate along this potential and gain kinetic energy; if this energy exceeds the trap depth U_0 , both atoms are lost from the trap in a single collision event. This process is illustrated schematically in Fig. 2.4.

² The ability to resolve multiple occupancies using fluorescence imaging has been demonstrated only in specialized configurations; see, for example, [doi:10.1103/PhysRevLett.111.253001](https://doi.org/10.1103/PhysRevLett.111.253001).




Figure 2.4: Molecular interaction potentials as a function of the internuclear separation. Two atoms in the ground S-state experience a van der Waals potential scaling as R^{-6} . When one atom is excited to the P-state by red-detuned light, the interaction becomes a resonant dipole-dipole potential scaling as R^{-3} . This leads to an acceleration of the atoms toward each other and their subsequent loss once their kinetic energy exceeds the trap depth.

Because each collision event removes the atoms in pairs, after a sufficient number of collisions each tweezer will be left with either zero or one atom, depending on whether the initial number of atoms was even or odd. The loss process therefore naturally saturates at a *binary occupation*, as no further pairwise collisions can occur once only one atom remains.

Assuming the initial atom number in each tweezer is Poisson-distributed with mean $\tilde{N}$, the probability of ending with a single atom corresponds to the probability that the initial occupation was odd. This yields an approximate single-atom loading probability

$$p_{\text{single}} \approx \frac{1}{2} \left(1 - e^{-2\tilde{N}} \right). \quad (2.49)$$

For large $\tilde{N}$, even and odd occupations become equally likely, and for $\tilde{N} \gtrsim 5$ the difference is already negligible. Consequently, on average about 50% of the tweezers end up empty, while the remaining 50% contain exactly one atom. It is worth noting that several techniques have been developed to overcome this probabilistic limit and achieve substantially higher single-atom loading efficiencies, including entropy redistribution, tailored loading protocols, and engineered dissipation. Such methods have been successfully demonstrated in alkali and alkaline-earth-like species such as ytterbium, enabling near-deterministic loading

of large tweezer arrays [64]. However, the application of these advanced loading schemes to strontium optical tweezer platforms has not yet been experimentally demonstrated.

2.1.7 Ground state cooling via Sisyphus technique

For single atoms confined in optical dipole traps or optical tweezers, the dynamics differ fundamentally from those in a magneto-optical trap. In optical traps, atoms experience a conservative optical potential that depends on the local light intensity and polarization, rather than a dissipative radiation-pressure force. As a consequence, optical tweezers provide confinement, but do not intrinsically cool the atomic motion.

To reduce the atomic temperature within such conservative potentials, additional cooling mechanisms must be employed. These mechanisms differ from Doppler cooling and are not fundamentally limited by the Doppler temperature that characterizes MOT operation. At sufficiently low temperatures, typically in the microkelvin regime, the atomic motion in the trap becomes quantized, and the dynamics are more accurately described in terms of discrete vibrational levels of the trapping potential.

Reaching the motional ground state of the trap is therefore a key requirement. Ground-state cooling minimizes the spatial extent and momentum uncertainty of the atomic wavefunction, which is crucial for coherent Rydberg-mediated interactions and for preserving confinement during fluorescence imaging. In imaging protocols, repeated photon scattering can otherwise induce heating and lead to atom loss, making tight motional control essential for high-fidelity detection and long experimental sequences.

Among the cooling methods suitable for optical tweezers, the two most relevant are *sideband cooling* [65, 66, 67] and *Sisyphus cooling* [68, 69, 70, 71]. Sideband cooling can be implemented in a magic trap, where the optical potentials for the ground and excited states connected by imaging transitions, U_g and U_e , are identical. In contrast, Sisyphus cooling operates in non-magic traps, where $U_g \neq U_e$, and remains effective whether the excited state is more tightly confined than the ground state ($U_e > U_g$) or vice versa ($U_e < U_g$).

In our experiment, the optical tweezers at 813.4 nm create a non-magic potential for the $^1S_0 \rightarrow ^3P_1$ transition. We implement *attractive Sisyphus cooling* using circularly polarized light to address the $m_J = 1$ sublevels. The basic principle of attractive Sisyphus cooling is illustrated schematically in Fig. 2.5. Classically, the atom oscillates around the trap

center, and the cooling laser is tuned to be resonant only at the bottom of the trap, where the detuning compensates for the local differential light shift $U_e - U_g$.

As the atom moves within the trap, it can absorb a photon near the potential minimum and be excited to the upper potential U_e , which is more tightly confining. The atom then climbs the potential hill toward the trap edges, increasing its potential energy. After a time determined by the excited-state lifetime, the atom spontaneously decays back to the ground state U_g , emitting a photon with higher energy than the absorbed one. The difference in photon energies corresponds to a net loss of kinetic energy, thereby cooling the atom. Repetition of many such cycles reduces the atom's motional energy, approaching the trap's ground-state energy.

A similar mechanism occurs in *repulsive Sisyphus cooling*, where $U_e < U_g$. In this case, the atom absorbs a photon when it is displaced from the trap center and loses energy while descending the excited-state potential before being optically pumped back to the lower potential. Even within the harmonic approximation, non-magic confinement leads to different motional eigenstates for $|g\rangle$ and $|e\rangle$.

Cooling of all motional degrees of freedom with a single beam is enabled provided that the cooling wavevector has non-zero projection onto each trap axis. [68, 70].




Figure 2.5: Schematic representation of the optical potentials for the ground (U_g) and excited (U_e) states in a non-magic optical trap. The atom absorbs a photon near the trap minimum, climbs the steeper excited-state potential, and then spontaneously decays to the ground state, losing kinetic energy in each cycle.

2.2 CHOOSING THE ELEMENT: STRONTIUM 88

The atomic structure of Strontium offers a unique combination of transitions with widely varying natural linewidths, ranging from broad dipole-allowed lines to ultranarrow intercombination and clock transitions. This feature enables highly efficient cooling at different experimental stages, precise coherent manipulation, and long coherence times.

Furthermore, the presence of both spin-singlet and triplet manifolds provides multiple excitation pathways, including access to Rydberg states, which are essential for engineering strong and controllable interactions between atoms.

This section will present an overview of the atomic properties of Strontium that are most relevant to our experiment. It begins with a description of its electronic level structure and the key transitions employed for laser cooling and detection, followed by a discussion of the ultranarrow clock transition and the excitation schemes used to reach Rydberg states.

2.2.1 *Electronic structure and level diagram*

Strontium is an alkaline-earth element with electron configuration $[\text{Kr}]5s^2$, that means that it is an atom with two valence electrons, a feature shared by all alkaline-earth elements (Be, Mg, Ca, Sr, Ba, Ra). This configuration, however, is not unique to this group: helium represents the simplest two-electron system, while ytterbium is another atom with two valence electrons that has been widely used in recent years for experiments closely related to those presented in this work [72].

One of the most relevant properties of strontium for our purposes is its level structure and the set of electronic transitions it offers, illustrated in simplified form in Fig. 2.6.

The two-valence-electron structure gives rise to a rich spectrum of singlet and triplet states. Particularly relevant to this work are the ground state $^1S_0$, the excited states $^3P_0$, $^3P_1$, and $^1P_1$, as well as Rydberg levels in which one electron is promoted to a state with high principal quantum number n . The $^1S_0 \leftrightarrow ^3P_0$ "clock" transition is ultra-narrow and underpins both optical atomic clocks and qubit implementations.

Strontium occurs naturally in four stable isotopes, summarized in Tab. 2. For our experiment, we choose ^{88}Sr due to its relative abundance and bosonic nature, corresponding to a zero nuclear spin. This choice eliminates complications arising from hyperfine structure and reduces sensitivity to magnetic fields.




Figure 2.6: Simplified energy level scheme for ^{88}Sr , highlighting transitions of interest.

The high natural abundance of ^{88}Sr enables efficient atom loading from a thermal beam, while the absence of nuclear spin simplifies laser cooling, detection schemes, and the implementation of Rydberg excitation. Nonetheless, the spectral shifts among the four isotopes are small enough to be addressed using standard acousto-optic modulators (AOMs), allowing, in principle, the use of all four isotopes with a single apparatus and laser system [73].

Table 2: Properties of the stable isotopes of Strontium [74].

Isotope	Natural abundance	Statistics	Nuclear spin
^{84}Sr	0.56%	bosonic	0
^{86}Sr	9.86%	bosonic	0
^{87}Sr	7%	fermionic	9/2
^{88}Sr	82.58%	bosonic	0

2.2.2 Key Transitions of Strontium

BLUE TRANSITION. The blue transition $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^1P_1$ has a wavelength of $\lambda = 460.9$ nm and a natural linewidth of $\Gamma/2\pi = 32$ MHz, corresponding to a lifetime of $\tau \sim 5.3$ ns [75]. This broad transition is ideal for the initial deceleration of the atomic beam in the Zeeman slower and in the 2D MOT, where it efficiently reduces the kinetic energy of the atoms before entering the main trapping region. It is also used for the first stage of the 3D MOT, providing strong cooling and confinement forces, as well as for absorption imaging of the MOT and fluorescence detection of single atoms in optical tweezers. The excited $^1P_1$ state can decay, with a small branching ratio, into the $5s4d\ ^1D_2$ state. This intermediate level subsequently decays into the metastable $5s5p\ ^3P_1$ and $^3P_2$ states, removing atoms from the cooling cycle and limiting the efficiency of the blue MOT.

RED TRANSITION. The red transition $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^3P_1$ has a wavelength of $\lambda = 689.4$ nm and a natural linewidth of $\Gamma/2\pi = 7.6$ kHz, corresponding to a lifetime of $\tau = 21$ μ s [76]. As a narrow intercombination line, it enables cooling to Doppler temperatures roughly three orders of magnitude lower than those achievable with the blue transition. Its narrow linewidth leads to a small capture velocity, which prevents direct trapping of a hot atomic beam in a red MOT, but the linewidth can be artificially broadened to increase the capture range, allowing for second-stage MOT operation and preparation of colder and denser atomic clouds. This transition also plays a central role in in-tweezer cooling (see Sec. 2.1.6) and high-resolution imaging of single atoms [55].

REPUMPER TRANSITIONS. The dipole-allowed transitions $5s5p\ ^3P_J \leftrightarrow 5s6s\ ^3S_1$ are essential for managing atomic populations within the ^3P_J manifold. The relevant wavelengths and parameters are:

- $^3P_0 \leftrightarrow ^3S_1$: $\lambda = 679.1$ nm, $\Gamma/2\pi = 1.4$ MHz, $\tau = 112$ ns
- $^3P_1 \leftrightarrow ^3S_1$: $\lambda = 687.8$ nm, $\Gamma/2\pi = 4.3$ MHz, $\tau = 37$ ns
- $^3P_2 \leftrightarrow ^3S_1$: $\lambda = 707.0$ nm, $\Gamma/2\pi = 6.7$ MHz, $\tau = 24$ ns

These transitions repump atoms that fall into long-lived metastable states as a result of decay from the $^1P_1$ state during cooling or imaging. Atoms in $^3P_1$ decay back to the ground state within a few microseconds, whereas atoms in the metastable $^3P_2$ remain trapped in a dark state.

A laser resonant with the 707 nm transition clears the $^3P_2$ population, and a second repump laser at 679 nm depletes the $^3P_0$ metastable population, ensuring efficient recycling of atoms into the cooling cycle. A dedicated repumper addressing the $^3P_1$ state at 688 nm is not strictly necessary, as this level is short-lived and rapidly decays to the ground state without external intervention. As a result, population accumulation in $^3P_1$ remains negligible under typical operating conditions of blue MOT, both in the standard and in the *red-shielded* configuration.

CLOCK TRANSITION. The clock transition $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^3P_0$ occurs at a wavelength of $\lambda = 698.4$ nm and is characterized by an extremely narrow natural linewidth, effectively $\Gamma \rightarrow 0$, in the bosonic isotopes of Sr, corresponding to a lifetime $\tau \rightarrow \infty$. In practice, the lifetime is limited only by a very weak two-photon E1–M1 decay channel to the $^1S_0$ ground state, with a theoretical lifetime on the order of 10^3 years[77]. The extreme narrowness arises because the transition is doubly forbidden by selection rules: $\Delta S = 0$ and $J = 0 \not\rightarrow J' = 1$. While the first rule is partially relaxed by fine-structure mixing of singlet and triplet states, the absence of hyperfine interaction in the bosonic isotopes of Sr (nuclear spin $I = 0$) enforces the second rule strictly.

Coherent excitation of this transition is possible by applying a strong external magnetic field ($B \sim 50$ G), which mixes the ^3P_J states and opens the transition channel (see Sec. 3.1.2). Under these conditions, the effective linewidth scales as $\Gamma(B) \propto B^2$, while the Rabi frequency scales as $\Omega(B) \propto B\sqrt{I}$, where I is the intensity of the driving laser[78]. In the fermionic isotope ^{87}Sr , with nuclear spin $I = 9/2$, the hyperfine interaction naturally mixes electronic states, resulting in a small but nonzero natural linewidth without the need for an external magnetic field. The metastable $^3P_0$ state plays a central role in optical atomic clocks and qubit implementations, and it can also serve as an intermediate level for coherent excitation to Rydberg states.

RYDBERG EXCITATION TRANSITIONS. The transitions $5s5p\ ^3P_0 \leftrightarrow 5sns\ ^3S_1$ (at $\lambda \sim 317$ nm) and $5s5p\ ^3P_1 \leftrightarrow 5sns\ ^3S_1$ (at $\lambda \sim 319$ nm) provide a convenient access to Rydberg states with $n \gg 1$, which are essential for enabling strong, long-range interactions between otherwise non-interacting neutral atoms. Each of these Rydberg series spans a spectral range of approximately 1.5 nm, extending from low principal quantum numbers up to the ionization threshold.

2.3 TUNING THE INTERACTION: RYDBERG ATOMS

Arrays of individually trapped neutral atoms in optical tweezers provide a scalable and highly controllable platform for quantum simulation. In their ground state, however, the atoms are separated by several micrometers and do not naturally interact on relevant timescales, which strongly limits the range of many-body phenomena that can be explored. To overcome this limitation, strong and controllable long-range interactions can be introduced via excitation to Rydberg states.

Rydberg states are highly excited electronic states with large principal quantum number n , where one electron occupies a very extended orbital. Their large spatial extent makes them an ideal resource for engineering long-range atomic interactions and probing correlated quantum phenomena.

2.3.1 Rydberg physics: scaling laws and blockade

Rydberg atoms are defined as atoms in which a valence electron is promoted to a state with very high principal quantum number n . In this regime, the atom resembles a hydrogenic system, with a single electron moving far from the nuclear core. The influence of the inner electrons on the Rydberg electron is not negligible and is described through the quantum defect $\delta_{n,\ell,j}$, which depends on n , the orbital angular momentum ℓ , and the total angular momentum j . The binding energy of a Rydberg state can therefore be written as [79]:

$$E_{\text{bind}} = \frac{Ry}{(n - \delta_{n,\ell,j})^2}, \quad (2.50)$$

where Ry is the Rydberg constant corrected for the reduced mass of the electron–nucleus system:

$$Ry = \frac{m_e}{M + m_e} \frac{m_e e^4}{8\epsilon_0^2 h^2}, \quad (2.51)$$

with M the atomic mass, m_e the electron mass, e the elementary charge, h Planck's constant, and ϵ_0 the vacuum permittivity. From Eq. 2.50 one sees that the binding energy scales as $\sim n^{-2}$, which implies that the level spacing between adjacent Rydberg states decreases as $\sim n^{-3}$. These scaling behaviors, as well as several other characteristic properties, have been extensively studied both theoretically and experimentally [79, 80, 81] and are reported in Tab. 3.

Table 3: Scaling of key Rydberg atom properties with principal quantum number.

Property	n -scaling
Binding energy	n^{-2}
Energy between adjacent n states	n^{-3}
Radiative lifetime τ	n^3
Orbital radius $\langle r \rangle$	n^2
Dipole moment $\langle np er nd \rangle$	n^2
Scalar polarizability α	n^7
Dipole-dipole coefficient C_3	n^4
van-der-Waals coefficient C_6	n^{11}

The spontaneous emission rate towards a lower state is given by:

$$\Gamma \propto \frac{|\mu|^2 \omega^3}{3\pi\epsilon_0 \hbar c}, \quad (2.52)$$

where μ is the transition dipole moment and ω is the transition frequency. Spontaneous decay into neighboring Rydberg states is strongly suppressed due to the small transition frequency ω , while decay into the ground state is inhibited by the very small dipole matrix element μ , which originates from the negligible overlap between the corresponding wavefunctions. However, at finite temperature, blackbody radiation (BBR) induces additional couplings between nearby Rydberg levels. As a result, the dominant decay channels for high- n Rydberg states are radiative decay to the ground state and BBR-induced transitions to neighboring Rydberg states with principal quantum numbers $n \pm 1, 2$. These processes play a crucial role in determining the effective lifetime and population dynamics of Rydberg states under typical experimental conditions [82, 83, 84].

The most interesting feature of Rydberg atoms in the context of quantum simulation is the strength of their mutual interactions. For two atoms separated by a distance larger than the typical size of the electronic wavefunction $R \gg \langle r \rangle$, the dominant term is the electric dipole-dipole interaction Hamiltonian [2]:

$$\hat{V}_{dd} \sim \frac{\hat{\mathbf{d}}_1 \cdot \hat{\mathbf{d}}_2}{4\pi\epsilon_0 R^3}, \quad (2.53)$$

where the dipole operator is defined as $\hat{\mathbf{d}} = -e\hat{\mathbf{r}}$ and its matrix element between states with the same n is $d \sim n^2 ea_0$.

Two important regimes of Rydberg–Rydberg interactions can be distinguished. When both atoms are excited to the same spherically symmetric state, such as $|nS\rangle$, the average dipole moment vanishes and the first-order contribution of the dipole–dipole interaction is zero. In this case, the interaction appears only at second order in perturbation theory, through the virtual coupling of the pair state $|nS, nS\rangle$ to opposite-parity states like $|nP, nP\rangle$, which are detuned by Förster energy defect Δ . The resulting potential takes the form of a van-der-Waals interaction:

$$V(R) \propto \frac{C_6}{R^6}, \quad C_6 \propto \frac{d^4}{\Delta} \propto n^{11}, \quad (2.54)$$

and can reach values on the order of megaHertz for atoms separated by only a few micrometers in high- n Rydberg states [85, 86].

If the two atoms are prepared in states of different parity, for instance one in $|nS\rangle$ and the other in $|nP\rangle$, the dipole–dipole term contributes already at first order in perturbation theory. In this case, the interaction is a dipolar interaction:

$$V(R) \propto \frac{C_3}{R^3}, \quad C_3 \propto d^2 \propto n^4. \quad (2.55)$$

The coexistence of van der Waals and resonant dipole–dipole interactions provides an exceptional degree of flexibility in neutral-atom experiments. In addition to selecting the Rydberg quantum numbers or adjusting the separation of the optical tweezers, the interaction strength and character can be actively tuned by applying external electric or magnetic fields. Such fields shift the energies of the relevant pair states via Stark or Zeeman effects, allowing the system to be brought into or out of a Förster resonance, where the interaction changes from a van der Waals to a resonant dipole–dipole form and is strongly enhanced [87, 88, 89]. This tunability is a crucial feature for implementing programmable quantum simulators and engineered many-body Hamiltonians.

The strong distance-dependent interactions between Rydberg atoms not only determine the energy shifts of pair states but also give rise to a collective phenomenon of fundamental importance: the Rydberg blockade. Consider the resonant laser excitation of atoms initially in the ground state $|g\rangle$ to a Rydberg state $|r\rangle$. When two atoms are sufficiently close, the van-der-Waals interaction between them shifts the doubly excited state $|rr\rangle$ out of resonance with the excitation laser. As a consequence, the simultaneous excitation of both atoms is prevented, and only one of them can be promoted to the Rydberg level (see Fig. 2.7b).




Figure 2.7: Schematic illustration of the Rydberg blockade mechanism. (a) A single atom undergoes resonant laser excitation coupling $|g\rangle$ to $|r\rangle$ with Rabi frequency Ω_R . (b) For two atoms separated by $R < R_b$, the symmetric singly excited state $|\psi_+\rangle = (|gr\rangle + |rg\rangle)/\sqrt{2}$ is accessible, while $|rr\rangle$ is shifted out of resonance. (c) In a larger ensemble, the excitation of one atom blocks all atoms within a blockade radius R_b , enforcing collective dynamics.

This effect defines a characteristic distance known as the blockade radius R_b , inside which the excitation of a single Rydberg atom blocks the excitation of its neighbors. It is convenient to define R_b as the distance at which the interaction energy equals the Rabi coupling of the laser:

$$\frac{C_6}{R_b^6} = \hbar\Omega_R \quad \Rightarrow \quad R_b = \left(\frac{C_6}{\hbar\Omega_R} \right)^{1/6}. \quad (2.56)$$

The blockade mechanism plays a central role in many proposals and realizations of quantum simulation and quantum information protocols. For instance, two atoms initially prepared in $|gg\rangle$ under global laser excitation cannot populate $|rr\rangle$, and the system evolves in the symmetric subspace spanned by $|gr\rangle$ and $|rg\rangle$. This results in the creation of a maximally entangled state:

$$|\psi_+\rangle = \frac{1}{\sqrt{2}}(|gr\rangle + |rg\rangle), \quad (2.57)$$

with an enhanced collective Rabi frequency $\sqrt{2}\Omega_R$ [90].

The same feature can be extended to ensembles of N atoms confined within a blockade sphere. In this case, at most one excitation is allowed, and the system evolves into a symmetric superposition of all possible single excitations:

$$|\psi_N\rangle = \frac{1}{\sqrt{N}} \sum_{i=1}^N |g \dots r_i \dots g\rangle, \quad (2.58)$$

driven at a collectively enhanced Rabi frequency $\sqrt{N} \Omega_R$. In arrays of optical tweezers, where both the trap separation and the choice of Rydberg level can be controlled with high precision, the blockade mechanism allows the engineering of tunable long-range correlations extending across several sites.

2.3.2 Optical excitation pathways

Direct excitation of Strontium atoms to Rydberg states with a single photon from the $^1S_0$ ground state is experimentally impractical, since it would require deep-ultraviolet laser sources. A more convenient strategy is to use multi-step excitation schemes. Two relevant approaches are illustrated in Fig. 2.8.




Figure 2.8: Two excitation schemes to access Rydberg states in Strontium. (a) A first step excites atoms to the $^3P_0$ state via the clock transition at 698 nm, followed by a second photon at 318 nm. (b) A two-photon transition via the $^3P_1$ intermediate state, using 689 nm and 316 nm light.

In the first scheme (Fig. 2.8a), atoms are transferred to the long-lived metastable clock state $5s5p\ ^3P_0$, which can be treated as an effective ground state. From there, a single ultraviolet photon at ~ 318 nm (n -dependent in a range 317.5-320 nm) couples directly to a Rydberg state. This method has the advantage of effectively realizing a single-photon excitation without the drawbacks of intermediate state decay, while enabling relatively large Rabi frequencies ($\Omega \approx 10$ MHz [6]).

An alternative approach (Fig. 2.8b) exploits a two-photon process via the $5s5p\ ^3P_1$ level, where the first laser at 689 nm couples off-resonantly to the intermediate state, and a second photon around 316 nm (n -dependent in a range 315.5-317 nm) completes the excitation to a Rydberg state. Despite the apparent complexity, this method can drive the transition with coherence comparable to a single-photon process, at MHz frequencies [91, 92], provided that the detuning from the intermediate state is sufficiently large.

These two excitation schemes are the most widely used in neutral Strontium experiments, although alternative pathways involving both triplet and singlet states have also been demonstrated [93, 94, 95]. Our UV laser, described in Sec. 3.1.3, can be tuned between 316 nm and 319 nm, allowing us to span the energy difference between the $^3P_0$ and $^3P_1$ states, making both these excitation schemes accessible.

ARCHITECTURE OF THE EXPERIMENTAL MACHINE

This chapter provides an overview of the architecture and core components of the experimental system used throughout this work. The apparatus is designed to enable the trapping, cooling, manipulation, and imaging of ultracold strontium atoms. At its foundation lies the vacuum system, which establishes the ultra-high-vacuum (UHV) environment necessary to suppress collisions with background gas. To produce and control cold atomic ensembles, the setup incorporates a magnetic-field system centered on a high-field pair of anti-Helmholtz coils that generate the magnetic-field gradient required for magneto-optical trapping. This primary pair is complemented by three orthogonal sets of compensation coils, whose role is to cancel ambient magnetic fields and thereby ensure a stable and well-controlled environment around the atoms.

The vacuum and magnetic subsystems are integrated with a comprehensive optical setup containing all laser systems required for the sequential stages of atom preparation and control, including slowing, cooling, trapping, coherent manipulation, and imaging.

For practical and organizational reasons, the apparatus is built on two optical tables: one dedicated to the laser systems (the *laser table*) and the other hosting the experiment (the *experimental table*), as shown in Fig. 3.1. The experimental table supports two main breadboards: one housing the vacuum system and the other containing the optics used for laser cooling and atomic manipulation. The vacuum assembly is mounted on a separate breadboard equipped with a precision translation stage, enabling the entire structure to be moved forward and backward with high accuracy. This mechanical isolation from the rest of the apparatus not only improves stability, but also significantly simplifies alignment procedures and routine maintenance.

3.1 HARDWARE: THE ATOM-LIGHT MACHINE




Figure 3.1: Side view of the experimental optical table. On the left is the vacuum system, mounted on a separate breadboard supported by a translation stage. The optical setups for trapping, manipulating, and detecting the atoms are assembled on three levels of the optical table.

3.1.1 *Vacuum system and atomic source*

The vacuum system, shown in Fig. 3.2, consists of three main sections: the atomic source, a pumping chamber, and the rectangular glass cell. The full assembly extends approximately 105 cm from the oven in the atomic source to the distal end of the science cell.

The atomic source is a commercial *Sr Beam RevC* unit from AOSense, a compact, integrated module that combines a strontium oven, permanent magnets for the Zeeman slower and the 2D MOT, several optical access points for diagnostics and spectroscopy, a differentially pumped angled tube, a miniature ion pump, and a gate valve that allows isolation from the remainder of the vacuum system. The oven is powered via two electrical feedthroughs and heats a resistive element to approximately 420°C (measured by producer's sensor on the sample), thereby producing a thermal vapor of strontium. The resulting atomic beam passes through a chamber, which pre-collimates it by selecting a narrow distribution of transverse velocities. Immediately downstream, three viewports allow spectroscopic measurements of the atomic beam (see Ref. [96] for details).




Figure 3.2: 3D render of the vacuum apparatus.




Figure 3.3: Sr Beam RevC unit from AOSense. The dashed line indicates the trajectory of undeflected atoms, while the solid line represents the trajectory of deflected atoms.

From the outside, the permanent magnet assembly appears as a 5 cm long block located between the oven and the 2D MOT region. The Zeeman-slower beam at 461 nm enters through the “Hot” window, a sapphire viewport heated to 300°C to prevent coating by strontium and loss of transmission.

The slowing beam counter-propagates against the atomic flux and reduces the atoms' longitudinal velocity via the spatially varying magnetic field (see Sec. 2.1).

In addition to Zeeman slowing, transverse cooling and deflection are provided by two sets of 2D MOT beams entering through large side viewports and retro-reflected by in-vacuum optics. These beams further collimate and redirect the atomic beam to match the angle of the differential pumping tube before the cold-beam output flange (see Sec. 2.1.2). Efficient coupling into this tube requires careful optimization of beam geometry, polarization, detuning, and power. The recommended AOSense parameters served as an effective starting point. During the first year of operation, we introduced minor adjustments to maximize the 3D MOT atom number and minimize temperature. The optimized parameters are listed in Table 4.

Table 4: Optimized beam waist, power, detuning, and polarization for the atomic-source beams.

Beam	Waist	Power	Detuning	Polarization
Zeeman slower	8 mm (round)	50 mW	−580 MHz	σ^-
2D MOT (1)	25 mm $\times$ 5 mm	16 mW	−40 MHz	σ^+
2D MOT (2)	25 mm $\times$ 5 mm	8 mW	−40 MHz	σ^+

Spectroscopy on the atomic beam, used for frequency-locking the 461 nm laser, is performed by sending a retro-reflected 461 nm beam orthogonally through the atomic flux via dedicated vertical viewports (“Beam Lock Windows”). The retro-reflection compensates for possible transverse Doppler shift due to imperfect orthogonality and doubles the absorption signal. The 461 nm laser is locked to the $^1S_0 \leftrightarrow ^1P_1$ transition in ^{88}Sr using a double-pass AOM for fine frequency tuning. The diode current is modulated at 250 kHz, and the resulting absorption signal is demodulated to generate the error signal used for stabilization (see Sec. 3.1.3). Owing to the large natural linewidth ($\Gamma/2\pi = 32$ MHz) and the beam geometry, no saturation dip is observed; instead, a single broad absorption profile is detected.

The Zeeman-slower beam is delivered through an optical fiber and collimated with a Schäfter–Kirchhoff collimator (model 60FC-L-4M50L-01), producing the recommended 8 mm waist. A slight longitudinal focus improves slowing efficiency. Its polarization is set using a half-wave plate followed by a quarter-wave plate positioned in front of the *Hot* window.

The 2D MOT beam requires a 5:1 elliptical profile. This is produced using cylindrical lenses that expand the horizontal axis. The beam is then split using a half-wave plate and a PBS to balance the optical power between the two arms. Each arm subsequently passes through a pair of waveplates to set the required polarizations. The complete optical layout of the Zeeman slower, 2D MOT, and spectroscopy beams is shown in Fig. 3.4.




Figure 3.4: Optical setup for the Zeeman slower, 2D MOT, and spectroscopy beams.

The remainder of the vacuum apparatus consists of a custom octagonal pumping intermediate UHV section and houses all components required to sustain and monitor chamber and the rectangular glass science cell (see Fig. 3.2). The pumping chamber serves as an ultra-high vacuum: a NEX-Torr D500 StarCell pump (combining an ion pump and a NEG cartridge), and a Bayard–Alpert gauge (UHV-24P, Agilent), which provides pressure readings from 10^{-3} down to 5×10^{-12} Torr. Achieving such pressures requires a multistage pumping and bakeout procedure, with each stage optimized for a specific pressure regime. Further details on the assembly and bakeout of the vacuum system are reported in Ref. [97].

At the core of the experimental region lies the fused silica rectangular science cell, manufactured by Japan Cell, with internal dimensions

of $25 \text{ mm} \times 33 \text{ mm} \times 120 \text{ mm}$ (for x, y, z in Fig. 3.2). The surfaces are optically polished to a flatness of $\lambda/10$, minimizing optical distortion of the laser beams. The two large faces, that are employed for focusing imaging and tweezers light through the objectives, are AR-coated externally for the 460–815 nm range at normal incidence ($\text{AOI} = 0^\circ$), while all remaining faces are left uncoated.

The cell is connected to the octagonal chamber via a glass-to-metal junction mounted on a CF40 flange, which allows for a small degree of rotational freedom about the longitudinal axis. To align the cell, a reference 1 mm waist blue beam and two fixed-height alignment targets were used. By slightly rotating the flange, the reflections from the inner and outer glass surfaces were made to overlap with the reference beam path, ensuring that the optical axis of the cell was parallel to the optical table.

3.1.2 *Magnetic field control*

Precise magnetic field control is essential for several stages of the experiment, despite the insensitivity of bosonic ^{88}Sr in its ground state. Magnetic fields are required to operate both blue and red MOTs, to shift the position of the narrow-line MOT, to drive the clock transition, and to tune atomic polarizabilities. To meet these requirements, we employ a dual coil system consisting of a pair of high-field coils and three orthogonal compensation coil pairs, all mounted around the science cell on a custom non-conductive structure.

The main design goals include:

- magnetic-field gradients exceeding 50 G cm^{-1} for MOT operation;
- the ability to generate fields above 900 G along one axis for clock-state excitation;
- large optical access through the glass cell;
- independent control along all three spatial directions;
- high field uniformity over the few-hundred-micrometer extent of the tweezer array;
- efficient thermal management to suppress coil heating and prevent drifts.

A pair of water-cooled hollow-core copper coils provides the high fields and gradients required for both Helmholtz and anti-Helmholtz modes. These coils are currently used in anti-Helmholtz configuration for generating the MOT quadrupole field; a dedicated H-bridge circuit is under development to allow fast switching between the two configurations. The full design procedure, COMSOL simulations, and detailed experimental characterization of these coils have already been presented in a previous thesis [98], and are only summarized here. Coils CAD rendering is shown in Fig. 3.5




Figure 3.5: CAD rendering of one of the two high field coils: three pairs of layers of windings are embedded in an epoxy resin structure that serves as mounting structure too. The water Swagelok inlets and outlets and the electrical connection plates are visible. The electrical connections between the windings are in series while the hydraulic connections are in parallel. The rhomboidal shape is the one that leaves most of the optical access clear while guaranteeing high and uniform fields (in Helmholtz configuration) and gradients (in anti-Helmholtz configuration) at the center of the glass cell.

Three additional orthogonal compensation-coil pairs provide low-field control for stray-field cancellation, MOT positioning, and defining the quantization axis. Each pair is independently driven by a low-noise current source in Helmholtz configuration, producing uniform fields of a few Gauss with millisecond response times. Their geometry and placement are optimized to preserve optical access while ensuring mechanical

stability for the main coil assembly. Together, the high-field coils and compensation coils form an integrated magnetic environment that combines large tunability, excellent spatial uniformity, and fast switching capability, fulfilling the magnetic-field requirements of all stages of the experiment.

The main (high-field) coils produce large, uniform fields in Helmholtz mode (about 988 G at 190 A, corresponding to ≈ 5.18 G/A), and strong gradients in anti-Helmholtz mode (measured at 190 A: 128, 96, and 224 G/cm along x, y, z respectively). The full coil circuit (including 20 m of cables) has a resistance of 75.7 m Ω , corresponding to about 2.7 kW dissipation at 190 A. A dedicated snubber/IGBT discharge circuit produces a field extinction time of approximately 300 μ s (the theoretical snubber constant is ≈ 60 μ s but practical parasitic capacitance slows the decay). The compensation coils provide low-field, high-bandwidth magnetic bias control, with efficiency of 4–7 G/A and electrical time constants of 12–14 ms. These coils are used for MOT positioning and bias corrections. For more detailed design information and COMSOL modeling, see Ref. [98].

3.1.3 *Laser system and stabilization*

This section will describe the laser systems employed in the experiment. The setup currently includes:

- a blue laser addressing the $^1S_0 \rightarrow ^1P_1$ transition at 461 nm;
- a red laser for the $^1S_0 \rightarrow ^3P_1$ intercombination line at 689 nm;
- a tweezer laser at 813 nm, far detuned from any atomic transition and used for optical dipole trapping;
- a UV laser tunable between 316 nm and 319 nm for excitation to Rydberg states;
- two repumper lasers at 679 nm and 707 nm used to depopulate the metastable $^3P_0$ and $^3P_2$ states, maintaining a closed cooling cycle during blue MOT operation.

In the near future, a laser addressing the clock transition will also be integrated into the system.

Blue 461 nm laser

The blue laser system addressing the $^1S_0 \rightarrow ^1P_1$ transition at 461 nm is based on an injection-locked amplifier (ILA) manufactured by MOGLabs. The seed source is a Nichia NDB4916T diode, tunable between 450 nm and 465 nm, delivering up to 250 mW in an external-cavity diode laser (ECDL) configuration. Approximately 30 mW seed the ILA, while the remaining ~ 200 mW are available from an auxiliary output port. The ECDL is typically operated at 290 mA, and the ILA diode at 700 mA, yielding more than 650 mW of output power delivered to the experiment via an optical fiber.

A schematic of the optical layout is shown in Fig. 3.6. Two main branches originate from the ECDL and the ILA output. Both beams are shaped to produce a near-Gaussian, circular profile with a ~ 1 mm waist. On the seed side, one branch is directed to a wavelength meter (High-Finesse WS8-10, 10 MHz absolute accuracy) for frequency monitoring. Another branch is sent to the main optical table for Doppler-free spectroscopy on the AOSense Strontium source, providing the error signal for frequency stabilization. A final seed branch generates the 2D MOT beams, delivering 30 mW after the optical fiber.

The ILA output supplies the high-power beams used for the Zeeman slower (~ 60 mW) and the 3D MOT (~ 20 mW), both fiber-coupled. Two further branches are used for imaging: one for time-of-flight absorption imaging of the blue and red MOTs, and one for fluorescence imaging of atoms in tweezers. The fluorescence beam is slightly red-detuned to compensate for the tweezer-induced light shift; both imaging beams deliver ~ 0.5 mW at the experiment.

Frequency stabilization is implemented by modulating the diode current at 250 kHz while scanning the external-cavity piezo. The spectroscopy signal is detected with a balanced photodiode and demodulated by the ECDL controller to obtain a dispersive error signal. A built-in PID loop stabilizes both current and piezo, with user-adjustable parameters.

To obtain the large detuning required for the Zeeman slower (around -580 MHz), the laser is locked approximately 230 MHz below resonance and subsequently shifted by a single-pass 350 MHz AOM (G&H 3350-199) operated in the -1 diffraction order. This solution avoids inefficient double-pass geometries and the need for very high-frequency AOMs. The spectroscopy and imaging branches employ 110 MHz double-pass AOMs (G&H 3110-120), while the 2D and 3D MOT beams use 200 MHz




Figure 3.6: Schematic layout of the optical setup for the 461 nm laser. The asterisk fiber coupling shows the ILA output connection, from laser case to optical downstream paths.

single-pass AOMs (G&H 3200-121), all operated in the +1 order. The corresponding driving frequencies are indicated in Fig. 3.6.

Red 689 nm laser

The red laser addressing the $^1S_0 \rightarrow ^3P_1$ transition at 689 nm is a Titanium:Sapphire laser (SolsTiS, MSquared) pumped by an Equinox system delivering up to 18 W at 532 nm. The Ti:Sa crystal provides tunable red to near-infrared output, with high-fidelity two-photon driving of the Rydberg-transition frequency selection $\sim$ MHz that will be used to achieved using intracavity birefringent filters and an etalon. A Ti:Sa was chosen for its low high-frequency phase noise, compatible with the Rabi frequencies used in the experiment [99].





Figure 3.7: Schematic layout of the optical setup for the 689 nm laser.

Under typical conditions, about 100 mW at 689 nm suffice for locking to an ultra-low-expansion (ULE) cavity, operating the red MOT, performing Sisyphus cooling in the tweezers, and monitoring the frequency on a wavelength meter. For two-photon Rydberg excitation (Sec. 2.3.2), higher power is required; at full pump power, the system delivers up to 1.3 W.

Figure 3.7 shows the optical layout. The output is split into several branches for the ULE lock, MOT, Sisyphus cooling, two-photon Rydberg excitation, and wavelength-meter monitoring. The ULE-lock, Sisyphus, and MOT beams are frequency shifted by a double-pass 80 MHz AOM (G&H 3080-125). The Rydberg excitation beam uses a single-pass 350 MHz AOM (G&H 3350-199). The wavelength-meter branch is unshifted.

Frequency stabilization is performed via a Pound–Drever–Hall (PDH) lock to a notched 100 mm-long ULE cavity (Stable Laser Systems). The cavity mirrors are highly reflective at 689 nm, 698 nm, and 633 nm, allowing stabilization of the red MOT laser, the clock laser, and the Rydberg laser respectively (Sec. 3.1.3). The cavity is maintained under UHV at 1×10^{-7} Torr and thermally shielded by nested in-vacuum enclosures with a thermal time constant exceeding 30 hours.

In Figure 3.8a we can see how f_2 frequency offset is adjusted to bridge the frequency difference between the atomic resonance (carrier frequency) and one of the ULE cavity modes (f_1 sideband frequency).



(a) Illustration of the frequency components generated for offset Pound-Drever-Hall (PDH) locking. A $f_2 = 643$ MHz modulation creates two offset sidebands, each with a $f_1 = 8$ MHz PDH modulation.

(b) Frequency noise PSD of the 689 nm laser locked to the ULE cavity. The noise stays below the Beta Line above 500 Hz, corresponding to a linewidth compatible with the 7.6 kHz red transition.

Figure 3.8: 689 nm laser frequency stabilization scheme using PDH locking: (a) offset-PDH scheme and (b) frequency noise spectrum.

The measured frequency-noise spectrum of the ULE-stabilized laser confirms that noise remains below the Beta line [100] above 500 Hz, corresponding to a sub-kHz linewidth, sufficient for addressing the 7.6 kHz-wide red transition. Further improvements will be required for clock-transition spectroscopy, where the dominant noise is already approaching electronics limits.

Tweezer 813 nm laser

The tweezer laser at 813.4 nm is a Sirah Matisse Ti:Sapphire laser pumped by 25 W of 532 nm light from a Spectra-Physics Millennia EV25S. Up to 7.5 W can be coupled into a photonic-crystal fiber (NKT LMA-PM-10), which delivers the tweezer light to the experiment. An AOM placed before the fiber provides power control and fast switching.

Long-term stability of the laser system requires both frequency and power stabilization. Frequency locking is implemented using a HighFinesse WS8-10 wavelength meter, which provides feedback to the Matisse laser controller, ensuring stable and reproducible addressing of the clock transition. Power stabilization is achieved by sampling a fraction of the light after the output fiber and using a PID feedback loop to control the RF power driving the acousto-optic modulator (AOM) placed before the fiber.

Rydberg UV laser

Rydberg excitation requires ultraviolet light in the 316–319 nm range, which we generate by frequency-quadrupling the output of a Vexlum VALO SF VECSEL. The laser is tunable between 1266 nm and 1276 nm and delivers up to 1.1 W of infrared power. Each *Second-Harmonic-Generation* stage reduces the wavelength tuning range by a factor of two, and the final *Fourth-Harmonic-Generation* output covers from 316.5 nm to 319 nm, enabling excitation from either the $^3P_0$ or $^3P_1$ state (Sec. 2.3.2).

The FHG module (Leos S.r.l.) employs Hänsch–Couillaud locking [101] to maintain resonance in both SHG cavities. A small portion of the intermediate 633 nm light is extracted for frequency stabilization to the ULE cavity (Sec. 2.2.2). The module typically produces ~ 110 mW of UV power. Because UV optics generally introduce several percent of loss per element, and because scattered UV light is both hazardous and detrimental to the experiment, the optical path is kept as short as possible and uses a minimal number of components. For this reason, both the VECSEL and the FHG module are mounted directly on the main optical table, avoiding the need for optical fibers.

An AOM provides fast switching and power control. The UV output is shaped into a light sheet to enhance the intensity at the atomic sample and is initially overlapped with the 689 nm beam for two-photon excitation.

Repump lasers

Two Toptica DLC Pro lasers at 679 nm and 707 nm repump atoms from the metastable $^3P_0$ and $^3P_2$ states during blue MOT operation (Sec. 2.2.2), thereby restoring the closed cooling cycle. Each laser delivers ~ 20 mW and passes through an AOM for fast switching and intensity control. After modulation, the beams are combined on a polarizing beamsplitter and coupled into a single optical fiber, which delivers approximately 5 mW of each wavelength to the main table. These repumpers are essential for maximizing the MOT atom number, increasing it by roughly one order of magnitude (Sec. 4.1.1). Both lasers are frequency stabilized to the WS8-10 wavelength meter.

3.1.4 Optical setup for 3D MOT trapping

The optical setups for the blue and red three-dimensional MOTs are built directly on the main optical table, which is organized into three vertical levels, as shown in Fig. 3.9.




Figure 3.9: Layout of the optical setups for the blue and red 3D MOTs. The table is arranged in three vertical layers: the bottom level (below the cell) for branch selection, the middle level (at the cell height) for trapping in the horizontal plane, and the top level (above the cell) for vertical retro-reflection.

Bottom level. The blue and red MOT beams each originate from a dedicated polarization-maintaining fiber. A small fraction of the power is first sent to a photodiode for monitoring and active stabilization. Each beam is then split into three branches (one for each spatial the blue and red paths, which are then directed vertically toward the upper levels axis) using half-wave plates ($\lambda/2$) and PBSs. Three dichroic mirrors (red mirrors in Fig. 3.9a) overlay via periscopes. The blue absorption-imaging beam for MOT diagnosis is also introduced here, monitored and stabilized, and sent upward through its own periscope.

Middle level. This level hosts the glass science cell where the experiment takes place. The two horizontal MOT beam pairs intersect at 130° (instead of the conventional 90°), a geometry imposed by the size of the high-NA microscope objective. Achromatic quarter-wave plates ($\lambda/4$) set the appropriate polarization before retro-reflection. Because the large incidence angle on the cell windows reduces the transmitted intensity by $\sim 16\%$, two long-focal-length lenses ($f = 2500$ mm) are placed before

the retro-reflecting mirrors to refocus the beams and compensate for this imbalance.

Top level. The periscope delivers the vertical MOT beams, which are separated by a dichroic mirror and retro-reflected independently. This configuration was functional for an imaging system that is no longer in use. Quarter-wave plates flip the polarization of the beams in double pass upon return, while a half-wave plate in the blue path provides fine polarization control.

For both MOT wavelengths, the incoming beams have a waist of $w_0 = 8$ mm at the output of their respective fibers. Power stabilization is performed for each wavelength via a PID controller acting on the AOM RF drive.

The absorption-imaging beam for MOT diagnosis exits a fiber with a $w_0 = 1$ mm waist, reaches the cell via a periscope, and is collected by a Basler Ace acA4024-29um camera mounted beneath the upper breadboard (not shown in the figure).

A magnetic field gradient of $50.0(1)$ G cm⁻¹, generated by the high-field coils described in Sec. 3.1.2, provides the trapping force required for the blue MOT, while red MOT requires a lower gradient of 2 G cm⁻¹.

3.2 QUANTUM PLATFORM: OPTICAL TWEEZER GENERATION

The ability to generate arrays of optical tweezers represents a key step toward scalable quantum simulation platforms. On the one hand, it requires precise control over the spatial structure of the trapping light; on the other, it demands reliable techniques to load, cool, and organize atoms within the traps.

Various techniques can be employed to divide and shape a single incoming beam into multiple diffraction-limited spots, each acting as an individual trap for a single atom. Among the different available approaches, two technologies have emerged as the standard tools in neutral-atom experiments: *acousto-optic deflectors* (AODs) and *spatial light modulators* (SLMs). Although both enable the formation of multi-site trapping landscapes starting from a single laser source, they rely on fundamentally different physical principles and offer complementary strengths. AODs provide rapid, frequency-controlled beam steering with microsecond response times, while SLMs allow highly flexible, yet slower, wavefront engineering through programmable phase modulation. Together, they constitute the core optical-engineering techniques for generating reconfigurable arrays of optical tweezers.

Creating such arrays involves shaping laser beams into complex configurations and then focusing them through a high-numerical-aperture objective to produce tightly localized trapping potentials. Achieving deterministic single-atom occupancy further requires controlling the stochastic loading process, implementing efficient laser cooling strategies, and mitigating the randomness of initial atom distribution through active rearrangement into defect-free arrays.

3.2.1 High-NA Objective

The high-numerical-aperture ($NA = 0.55$) microscope objective used in this setup is a custom device manufactured by Special Optics. It serves a dual function: focusing the 813.4 nm trapping light to generate the optical tweezer array, and collecting fluorescence emitted by single atoms. For this purpose, the objective has been optimized across several wavelengths (461 nm, 689 nm, 698 nm, and 813 nm). It operates in the diffraction-limited regime at 461 nm, 698 nm, and 813 nm, exhibiting no focal shift between 461 nm and 813 nm and only a small $\sim 4 \mu\text{m}$ shift at 689 nm.

The objective has a working distance of 14 mm and is corrected for a 3.5 mm fused-silica cell. Its measured transmissions are 92.6% at 461 nm,

94.8% at 689 nm, and 96.6% at 813 nm. The mechanical housing is made of Ultem, a low-reflectivity polymer with excellent dielectric properties, well suited for operation in magnetic fields of several hundred Gauss. A summary of the optical and mechanical properties is provided in Tab. 5.

Table 5: Properties of the high-NA microscope objective.

Property	Value
NA	0.56 @461nm, 0.55 @813nm
Effective focal length	24 mm @461nm, 24.5mm @813nm
Aperture	27 mm
Field of view	0.36mm @813nm, 0.24mm @461nm 0.25mm @689nm
Working distance	9 mm vacuum + 3.5 mm silica + 1.5 mm air
Diffraction-limited	461nm and 813nm
Transmission	92.6% @461nm, 94.8% @689nm, 96.6% @813nm

The objective is mounted on a 5-axis assembly combining a pitch–yaw platform (PY004/M, Thorlabs), a long-range linear translation stage (L-509.44AD00, PI), and a two-axis translating lens mount (CXY2, Thorlabs). This configuration provides full control of tip, tilt, and translation along all spatial directions. The 102 mm travel of the translation stage allows the objective to be completely retracted from the coil structure around the science cell, which is advantageous both during alignment and for future maintenance. A detailed account of the alignment procedure can be found in Ref. [97].

3.2.2 Acousto-Optic Deflectors (AODs)

Acousto-Optic Deflectors (AODs) are optical devices that exploit the interaction between light and acoustic waves in a crystal [102]. When a radio-frequency (RF) signal is applied to a piezoelectric transducer bonded to the crystal, it generates an acoustic wave that propagates through the medium. This acoustic wave produces a periodic modulation of the refractive index, effectively forming a dynamic diffraction grating. A laser beam traversing the crystal is thus diffracted, and the deflection angle depends directly on the acoustic frequency.

The deflection angle θ of the first-order diffracted beam is given by [103]:

$$\theta(\nu) = \frac{\lambda}{v} \nu, \quad (3.1)$$

where λ is the wavelength of the laser light, v is the speed of sound in the crystal, and ν is the RF frequency driving the acoustic wave. This linear relation enables precise electronic control of the beam position through modulation of the RF signal [56].

AODs combine robustness with fast response and high angular resolution. Because they contain no moving parts, they provide stable and reliable beam steering, supporting rapid scanning and flexible manipulation of tweezer arrays in modern quantum simulation experiments.

The AODs used in our experiment are manufactured by AA OptoElectronics (model DTSX-400-810, specified for $\lambda = 813$ nm). They provide a measured operational bandwidth of approximately ± 20 MHz around the central driving frequency $\nu_0 = 101$ MHz. Using Eq. (3.1), the angular separation between the zeroth- and first-order diffracted beams at the central frequency is $\theta_0 = 0.13$ rad $\approx 7.2^\circ$.

Driving the AOD with multiple RF tones simultaneously produces several first-order diffracted beams at distinct angles, enabling the generation of multi-site tweezer arrays. Conversely, applying a time-dependent RF frequency allows the deflected beam to be steered dynamically, which can be used either for rapid point-by-point scanning or, when the scan rate exceeds the relevant dynamical timescales, for synthesizing an effectively time-averaged light distribution. In addition to angular control, AODs permit modulation of beam intensity through the RF amplitude and support fast beam switching on sub-microsecond timescales.

Generating Optical Tweezer Arrays with AODs

Placing a scan lens at its focal length from the AOD transforms angular separations into spatial separations (Fig. 3.10, top). Each beam deflected by the AOD at a different angle is focused to a different point in the focal plane of the lens. After passing through this focal plane, the beams propagate parallel to each other, forming a one-dimensional array of spatially separated beams. A tube lens and the microscope objective then focus these beams into an array of optical tweezers at the focal plane of the objective, all propagating in parallel. This is the principle for realizing a 1D array of optical tweezers using a single AOD.

For example, if the AOD is driven with two RF frequencies ν_1 and ν_2 , two beams are deflected at angles θ_1 and θ_2 . After passing through a telescope with magnification $M = f_{\text{tube}}/f_{\text{scan}}$, the angular separation is demagnified by M . The distance d between the two tweezers focused by the objective with focal length f_{obj} is given by [58]:

$$d = 2f_{\text{obj}} \tan \left[\frac{\theta_1 - \theta_2}{2M} \right] = 2f_{\text{obj}} \tan \left[\frac{(\nu_1 - \nu_2)\lambda}{2v} \frac{1}{2M} \right]. \quad (3.2)$$

Two-dimensional square or rectangular arrays of tweezers can be generated using two orthogonal AODs (Fig. 3.10, bottom) to independently deflect along both axes. A $1:1$ relay telescope in a $4f$ configuration is placed between the two AODs to match the deflection planes, ensuring that angular deflections map precisely into spatial shifts at the object plane.

Driving the AODs with static RF frequencies generates fixed one- or two-dimensional arrays, where each frequency tone corresponds to a single trap location. Dynamically modulating the RF frequencies allows real-time control of trap positions, enabling the creation of moving optical tweezers.




Figure 3.10: Optical layout for generating arrays of optical tweezers using AODs. (Top) A single AOD creates a 1D array of laser beams. A scan lens and a tube lens relay angular separations to the back aperture of the objective, which focuses the beams into a 1D array of optical tweezers. (Bottom) Two orthogonal AODs (AOD_x and AOD_y) allow independent control along both axes, producing a 2D array at the focal plane of the objective.

3.2.3 Spatial Light Modulators (SLMs)

Spatial Light Modulators (SLMs) are widely used to modulate laser beams for the generation of optical tweezer arrays. In quantum optics and atomic physics applications, they are typically employed as *phase-only* devices, enabling precise control over the wavefront of the incident light in order to synthesize arbitrary intensity patterns at the focus of a high-NA objective.

Among the available technologies, Liquid Crystal on Silicon (LCOS) SLMs are particularly popular because they offer high spatial resolution, phase-only modulation, and comparatively low optical loss. Each pixel consists of a thin layer of liquid crystal whose molecular orientation can be adjusted by an applied voltage. This modifies the effective refractive index experienced by the incoming polarized light, thereby imprinting a programmable phase shift on the reflected wavefront. The resulting two-dimensional phase mask $\phi(x, y)$ determines the far-field optical intensity. When a lens is placed one focal length downstream of the SLM, the imposed phase modulation is mapped into a spatial distribution of light in the focal (Fourier) plane, allowing the creation of complex tweezer geometries [104].

Fourier Optics Description

Within the paraxial approximation, the optical field incident on the SLM can be written as

$$U(x, y) = A(x, y)e^{i\phi(x, y)}, \quad (3.3)$$

where $A(x, y)$ denotes the amplitude of the incoming beam and $\phi(x, y)$ is the programmed phase mask. In the focal plane of a lens or microscope objective, the resulting field distribution is proportional to the two-dimensional Fourier transform of the modulated wavefront:

$$\tilde{U}(k_x, k_y) = \mathcal{F}[U(x, y)], \quad (3.4)$$

with (k_x, k_y) the spatial-frequency coordinates. By adjusting the pixel voltages on the SLM, target far-field intensity profiles $I_{\text{target}}(k_x, k_y)$ can be synthesized, including periodic or arbitrary two-dimensional arrays of optical tweezers [105].

Initial Phase Correction and Surface–Aberration Compensation

In any practical implementation, a spatial light modulator does not provide a perfectly uniform phase response across its surface. Manufacturing tolerances, pixel-to-pixel variations, and residual curvature of the LCOS backplane introduce spatially dependent phase offsets, which degrade the quality of the generated intensity pattern. Furthermore, the optical system downstream of the SLM—including relay optics and the high-numerical-aperture objective—may introduce additional aberrations such as astigmatism, coma, or spherical distortion. For high-fidelity tweezer generation, these static distortions must be compensated through an initial calibration phase before running any holographic optimization algorithm.

DEVICE- AND WAVELENGTH-DEPENDENT PHASE CALIBRATION. The phase shift imparted by each pixel of the SLM depends on the applied voltage and on the operating wavelength. As a result, the mapping between mask digital gray-levels and optical phase must be measured for each device and wavelength used in the experiment. This calibration is typically obtained by recording the interference between a reference beam and the SLM-modulated beam, allowing one to extract the effective phase response curve and correct for nonlinearity. A second, spatially resolved measurement provides a 2D “flatness map,” representing the residual phase profile of the device when nominally set to a constant gray level. This map, often slowly varying across the aperture, accounts for backplane bowing, pixel-fill-factor effects, and minor alignment offsets.

The resulting phase offset map $\phi_{\text{SLM}}^{\text{flat}}(x, y)$ is subtracted from all subsequent holograms:

$$\phi_{\text{corr}}(x, y) = \phi_{\text{target}}(x, y) - \phi_{\text{SLM}}^{\text{flat}}(x, y), \quad (3.5)$$

ensuring that the effective phase sent to the optical system is the one prescribed by the hologram-design algorithm.

ZERNIKE-BASED OPTICAL ABERRATION COMPENSATION. Beyond the SLM surface itself, static aberrations in the relay optics and microscope objective distort the wavefront and broaden or distort the generated traps. These aberrations can be characterized experimentally using either the optical response of a single trap or a more global metric such as the Strehl ratio across the field of view. [106] Once measured, they are most effectively represented using a basis of Zernike polynomials, which

provide an orthogonal basis for the expansion of low-order aberrations over a circular aperture:

$$\phi_{\text{Zernike}}(x, y) = \sum_{n=1}^N a_n Z_n(x, y). \quad (3.6)$$

The coefficients a_n are determined either through direct wavefront measurement (e.g. Shack–Hartmann sensing [107]) or via iterative optimization using the trap response as a feedback signal. Applying the negative of the recovered aberration to the SLM

$$\phi_{\text{opt}}(x, y) = -\phi_{\text{Zernike}}(x, y) \quad (3.7)$$

allows the SLM to pre-compensate for distortions in the optical path, effectively restoring a diffraction-limited focus at the trapping plane.

COMBINED CORRECTION PIPELINE. In practice, both types of corrections are applied simultaneously. The total calibrated phase mask written to the SLM is the sum of the holographic phase and all correction layers:

$$\phi_{\text{SLM}}(x, y) = \phi_{\text{target}}(x, y) - \phi_{\text{SLM}}^{\text{flat}}(x, y) - \phi_{\text{Zernike}}(x, y). \quad (3.8)$$

This ensures that the hologram operates on an effectively idealized, flat, and aberration-free phase surface. The resulting improvement in trap fidelity, uniformity, and stability is especially important for large-scale tweezer arrays and for experiments that require high phase stability such as rearrangement, transport, or coherent Rydberg manipulation.

Because the complete calibration procedure is highly device- and wavelength-specific, and is part of ongoing work within the host group, the detailed experimental characterization of these corrections is omitted here. Nonetheless, the framework described above constitutes an essential preprocessing stage for all hologram-generation routines used in this experiment

Phase Mask Design: Gerchberg–Saxton Algorithm

Because only intensity can be directly measured in the target plane, computing a phase mask that reproduces a desired intensity pattern requires an iterative phase-retrieval procedure. The most commonly used method in holographic tweezer generation is the Gerchberg–Saxton (GS) algorithm [108], which iteratively refines the SLM phase pattern until the propagated field closely matches the prescribed target intensity.

Starting from an initial guess for the SLM phase pattern $\phi_0(x, y)$, applied on the input amplitude $A(x, y)$

$$U_0(x, y) = A(x, y) e^{i\phi_0(x, y)}, \quad (3.9)$$

the complex field in the target (Fourier) plane is obtained via

$$\tilde{U}_0(k_x, k_y) = \mathcal{F}[U_0(x, y)]. \quad (3.10)$$

The amplitude is then replaced with the desired target intensity $I_{\text{target}}(k_x, k_y)$ while retaining the computed phase:

$$\tilde{U}'_0(k_x, k_y) = \sqrt{I_{\text{target}}(k_x, k_y)} e^{i \arg \tilde{U}_0(k_x, k_y)}. \quad (3.11)$$

Propagation back to the SLM plane gives

$$U_1(x, y) = \mathcal{F}^{-1}[\tilde{U}'_0(k_x, k_y)]. \quad (3.12)$$

The known input amplitude profile is then enforced, keeping the updated phase:

$$U'_1(x, y) = A(x, y) e^{i \arg U_1(x, y)}. \quad (3.13)$$

This sequence is iterated until convergence, typically assessed by monitoring the error between the generated and target intensities. Owing to its numerical stability, low computational cost, and straightforward implementation, the GS algorithm is a standard tool for holographic beam shaping in optical tweezer experiments.

Weighted and Neighborhood-Weighted implementations

Although the standard Gerchberg–Saxton (GS) algorithm is widely used for designing phase holograms for optical tweezer arrays, its convergence often degrades when the target pattern contains many traps with tight spacing or large dynamic range. To address this, two improved variants have been developed: the *Weighted GS* (WGS) algorithm and the *Neighborhood Weighted GS* (NeWGS) algorithm. [109, 110]

Figure 3.11 summarizes the common hologram-generation workflow used in their control software, including SLM correction maps and phase optimization loops.




Figure 3.11: Flowchart of the hologram-generation sequence implemented in the control software.

WEIGHTED GERCHBERG-SAXTON (WGS). WGS improves trap uniformity by introducing a dynamic weight vector w_j that compensates over- or under-filled traps at each iteration. In our implementation, the trap fields are computed individually in the SLM plane, summed coherently, and transformed to the Fourier plane as in the standard GS loop. After each iteration, trap intensities are evaluated and weights updated according to a multiplicative rule:

$$w_j^{(n+1)} = w_j^{(n)} \times \frac{I_{\text{target},j}}{I_{\text{meas},j}}. \quad (3.14)$$

Below is the pseudocode (Alg. 1) corresponding to the `WGS()` implementation.

NEIGHBORHOOD WEIGHTED GS (NeWGS). NeWGS extends the WGS method by modifying the trap weights using not only the trap's own error but also the errors of neighboring traps. This reduces local fluctuations, smooths the update landscape, and accelerates convergence for dense arrays.

Algorithm 1: Weighted Gerchberg–Saxton (WGS)

Input: Target intensities I_{target} , input beam $A(x, y)$
Output: Phase mask $\phi_{\text{SLM}}(x, y)$

- 1 Initialize phase $\phi_0(x, y)$ and weights $w_j = 1$
- 2 **for** $n = 1 \dots N_{\text{max}}$ **do**
- 3 Compute SLM-plane field $U_n = Ae^{i\phi_n}$
- 4 Propagate to Fourier plane $\tilde{U}_n = \mathcal{F}[U_n]$
- 5 Measure trap intensities $I_{\text{meas},j}$ at trap positions
- 6 Update weights: $w_j \leftarrow w_j \cdot (I_{\text{target},j} / I_{\text{meas},j})$
- 7 Construct weighted target field $\tilde{U}'_n = \sqrt{w_j I_{\text{target}}} e^{i \arg \tilde{U}_n}$
- 8 Propagate back: $U'_n = \mathcal{F}^{-1}[\tilde{U}'_n]$
- 9 Update phase $\phi_{n+1} = \arg(U'_n)$
- 10 **return** ϕ_{SLM}

Algorithm 2: Neighborhood Weighted GS (NeWGS)

Input: Target intensities I_{target} , trap positions, beam profile
Output: Phase mask ϕ_{SLM}

- 1 Initialize $\phi_0(x, y)$, weights $w_j = 1$, and neighborhood kernel K_{jk}
- 2 **for** $n = 1 \dots N_{\text{max}}$ **do**
- 3 Compute Fourier-plane field and trap intensities $I_{\text{meas},j}$
- 4 **foreach** trap j **do**
- 5 Compute neighborhood-weighted numerator and denominator:
- 6 $N_j = \sum_{k \in \mathcal{N}(j)} K_{jk} I_{\text{target},k}$
- 7 $D_j = \sum_{k \in \mathcal{N}(j)} K_{jk} I_{\text{meas},k}$
- 8 Update weight $w_j \leftarrow w_j \cdot (N_j / D_j)$
- 9 Construct corrected amplitude in Fourier plane
- 10 $\tilde{U}'_n = \sqrt{w_j I_{\text{target}}} e^{i \arg \tilde{U}_n}$
- 10 Propagate back and update phase $\phi_{n+1} = \arg \mathcal{F}^{-1}[\tilde{U}'_n]$
- 11 **return** ϕ_{SLM}

For each trap j , the weight update follows a neighborhood convolution:

$$w_j^{(n+1)} = w_j^{(n)} \times \frac{\sum_{k \in \mathcal{N}(j)} K_{jk} I_{\text{target},k}}{\sum_{k \in \mathcal{N}(j)} K_{jk} I_{\text{meas},k}},$$

where K_{jk} is a Gaussian-like kernel, and $\mathcal{N}(j)$ denotes the set of neighbor traps within a chosen radius.

The pseudocode above (Alg. 2) summarizes the procedure implemented in the functions `WGS()` and `NeWGS()`.

The WGS and NeWGS methods offer several improvements over standard GS:

- faster convergence for large or closely spaced tweezer arrays,
- improved trap-by-trap intensity uniformity (compensation of dynamic range),
- reduced stagnation and speckle artifacts,
- robustness to aberrations when combined with SLM non-ideality and Zernike corrections,
- compatibility with phase-locking for stable dynamical holograms.

A quantitative characterization of WGS and NeWGS - including trap homogeneity statistics, convergence studies, noise sensitivity, and comparison with standard GS - was carried out during my research visit at the *Institut d'Optique* on a Hamamatsu X15213 LCOS-SLM. Both algorithms were fully integrated into the hologram-generation software suite, based on the *python* classes `WGS()` and `NeWGS()` that extend the original GerchbergSaxton algorithm.

Because these results are part of ongoing research by the hosting group and they fall within the scope of intellectual protection agreements, they cannot be presented here in detail and with an exhaustive characterization.

In practice, the neighborhood-weighted update leads to a significantly faster and more stable convergence compared to standard WGS. For dense arrays, NeWGS typically reduces the number of iterations required to reach a given homogeneity by a factor of 2–3. Moreover, the final trap-to-trap intensity non-uniformity is commonly reduced from the few-percent level achieved with WGS (3–6 %) to the 1–3 % range. This improvement originates from the spatial averaging introduced in the weight-update rule, which suppresses local fluctuations and smooths the optimization landscape.

SLM Devices in Our Experiment

We plan to implement an SLM in our setup (Sec. 3.2.4). Two Meadowlark Optics devices are currently being tested:

- **E-19x12-500-1200-HDM8**: Resolution 1920×1200 pixels, pixel pitch $8 \mu\text{m}$, array size $15.36 \times 9.6 \text{ mm}$, HDMI controller with 60 Hz refresh rate and output trigger.
- **UHP SLM 1024x1024**: Resolution 1024×1024 pixels, pixel pitch $17 \mu\text{m}$, array size $17.40 \times 17.40 \text{ mm}$, analog data addressing, refresh rate 1.4 kHz.

These devices allow precise shaping of optical potentials and the creation of arbitrary 2D optical tweezer arrays for neutral-atom experiments.

3.2.4 *Tweezer generation setup*

The results presented in this thesis were obtained using two-dimensional arrays of optical tweezers generated via two crossed AODs, producing a static array. In the near future, the setup will be upgraded to combine an SLM for static array generation with the two crossed AODs, enabling dynamic rearrangement of atoms among the tweezers. A schematic of the final setup is shown in Fig. 3.12.

Laser Paths and Power Stabilization

A laser beam at 813 nm (Sec. 3.1.3) is split into two paths: one directed to the SLM and the other to the two crossed AODs. The two beams are recombined using waveplates and a polarizing beam splitter (PBS) in a collimated section of the beam path to minimize distortions. Both paths share a common telescope before the objective. Achromatic 2 inch lenses are used throughout the setup to reduce aberrations.

At the fiber output (60FC-SMA-T-4-M20-10 collimator), the laser beam has a waist $w = 1.8 \text{ mm}$. A waveplate and PBS convert polarization fluctuations into power fluctuations, which are monitored by a photodiode (PD). The photodiode signal is fed to a PID feedback loop to stabilize the fiber-delivered power.




Figure 3.12: Sketch of the optical tweezer generation setup employing two crossed AODs and an SLM.

AOD Branch

In the AOD path, a telescope expands the beam waist by a factor of 3.75 ($f_1 = 60$ mm, $f_2 = 225$ mm) to match the $7.5 \text{ mm} \times 7.5 \text{ mm}$ active aperture of the AODs. The beam is then sequentially deflected in the x and y directions by two orthogonal AODs. A dedicated 1 : 1 relay telescope ($f_1 = f_2 = 100$ mm) is placed between the two AODs to image the acoustic interaction region of the first device onto that of the second, ensuring proper overlap of the deflection planes.

A second relay telescope (focal lengths $f = 150$ mm, $f = 300$ mm) recombines the AOD branch with the SLM path and accommodates two steering mirrors and a dichroic mirror before the objective. The scan and tube lenses further magnify the beam by a factor of 2 ($f_{\text{scan}} = 150$ mm, $f_{\text{tube}} = 300$ mm), resulting in a beam waist at the back focal plane of the objective of $w = 13.5$ mm, matching the objective aperture radius $R = 13.5$ mm. The maximum extent of the tweezer array, obtained from the RF bandwidth (Eq. 3.2), is approximately $150 \text{ } \mu\text{m}$.

SLM Branch

In the SLM path, the beam is expanded by a factor of 10 using two telescopes ($f = 100$ mm, $f = 200$ mm and $f = -100$ mm, $f = 500$ mm) to match the dimensions of the UHP SLM 1024×1024 display. This configuration maximizes the active display area while minimizing power losses from beam clipping.

The phase-modulated light is relayed to the back focal plane of the objective via an additional telescope ($f = 300$ mm, $f = 200$ mm) and the final common telescope, producing a magnified phase-modulated image of approximately 24 mm at the back focal plane. A Semrock dichroic mirror with 3 mm substrate is placed before the objective to transmit the tweezer light while simultaneously reflecting the collected blue fluorescence along a separate optical path (Sec. 4.2.3). The substrate thickness limits aberrations introduced by the mirror to below $\lambda/4$.

Beam Size and Aperture Matching

The magnifications in the setup are optimized to match the beam size to the apertures of the devices (AODs and SLM), achieving a balance between efficiency and power loss. If the beam waist at the back focal plane is smaller than the objective aperture, the effective numerical aperture is reduced, producing larger tweezers. Conversely, if the beam waist exceeds the aperture, part of the Gaussian profile is clipped. Numerical simulations [60] show that the optimal compromise is obtained when $w/R \approx 1$.

3.2.5 Atom Rearrangement and Defect-Free Arrays

The Light-Assisted Collisions (LACs) process discussed in Sec. 2.1.6 ensures that each optical tweezer is occupied by at most one atom, leading to a random loading pattern with a filling fraction of about 50%. Even with more sophisticated loading techniques that can raise the single-site filling fraction above 80% [111, 112], some defects inevitably remain.

To create defect-free arrays in a deterministic way, the initially loaded atoms must be rearranged using an additional movable optical tweezer superimposed onto the static array. This auxiliary tweezer can pick up an atom from an occupied site and transport it to an empty target site [40]. While traditional spatial light modulators (SLMs) are well suited for generating static tweezer arrays, their refresh rates — typically tens of

milliseconds — are too slow for dynamic rearrangement. In contrast, a pair of orthogonal AODs driven by time-dependent RF signals can produce a tightly focused, dynamically steerable tweezer that is deeper than the static traps and suitable for fast atom transport.

Looking ahead, the development of fast spatial light modulators and three-dimensional trapping geometries opens the possibility of extending this approach to fully 3D tweezer arrays. In such configurations, atom reordering could be performed directly and in parallel by dynamically updating the SLM phase mask, eliminating the need for auxiliary transport tools, such as the AODs, and providing a significant advantage as the system size grows beyond one hundred atoms. Moreover, three-dimensional trapping enabled by SLMs naturally supports multi-plane imaging, where different layers of the array can be concurrently or sequentially or selectively probed by shifting them into the imaging focal plane after the experiment has been completed. These capabilities represent an important step toward scalable, high-density neutral-atom quantum simulators operating in three dimensions.

In our setup, a Spectrum M4i.6631-x8 AWG generates both the static RF tones and the frequency sweeps required for the rearrangement procedure. The two output channels operate independently and offer a 400 MHz bandwidth, a 1.25 GS/s sampling rate, 16-bit resolution, and 2 GS of on-board memory for storing the pre-calculated waveforms. After generation, the signals pass through a 5–150 MHz bandpass filter, are amplified using Mini Circuits ZHL-5W+ RF amplifiers, and are then delivered to the AODs. Fixed attenuators placed before the amplifiers ensure safe operation by matching the required 2 W input power.

The optimal rearrangement sequence is computed from the fluorescence image of the initially loaded array. Several algorithms have been proposed in the literature to minimize the number of required moves and reduce the overall rearrangement time [113, 114], but we propose our novel approach in Sec. 3.3.8.

A typical single move takes about 1 ms to transport an atom with high fidelity. For large arrays exceeding a few hundred atoms, the total reordering time can approach the vacuum-limited atom lifetime. To address this, some groups employ cryogenic environments to extend lifetimes [46], while others implement faster, parallelized rearrangement schemes using high-speed SLMs. Notably, a modified Gerchberg–Saxton algorithm implemented on a Meadowlark XY-series SLM operating at 2 kHz enabled parallel reordering with near-unity filling for arrays of up to 30 atoms [114].

3.3 SOFTWARE: CONTROL SYSTEM DEVELOPMENT

A central element of any modern atomic-physics experiment is the ability to coordinate a large number of heterogeneous devices with sub-microsecond synchronization and full programmability. In our setup, all aspects of sequence generation, hardware timing, data acquisition, parameter management, and automated analysis are implemented using the *Labscript Suite* [115], an open-source *python* framework increasingly adopted in Atomic Molecular Optical (AMO) laboratories. Labscript provides a scalable and hardware-agnostic structure that allows one to integrate AOMs, DDS-based frequency sources, acousto-optic and electro-optic modulators, shutters, cameras, NI-DAQ modules, FPGA boards, digital pulse generators, piezo controllers, and various custom electronics into a single coherent control environment.

3.3.1 *Labscript Suite*

In principle, Labscript cycles can be run simultaneously on different interconnected computers (via a local network), with each one performing a specific task, such as executing experimental shot routines or analyzing data, through dedicated GUIs. This setup would benefit for distributed computational load across multiple CPUs, preventing a single workstation from becoming overwhelmed. Crucially, references to a shot file and related metadata are passed between Labscript components and their worker processes via a ZeroMQ message bus.¹ This design enables efficient, low-latency data sharing and coordination between separate programs and processes, whether they are running locally or on distributed machines, thereby decoupling workflow execution from data analysis and device control. However, we are currently running everything on a single central computer.

The suite is organized around four main graphical interfaces, each responsible for a specific layer of the experimental workflow (Fig. 3.13): device communication, parameter management, timing visualization, and live data analysis. In the following, each component is described in detail, together with the integration strategies used to connect our experimental hardware to the Labscript environment.

¹ Labscript uses ZeroMQ as a publisher-subscriber and request-reply message bus to transmit shot file references and status messages between components (BLACS, Runmanager, Scheduler) and their worker processes. This decoupled messaging layer supports distributed execution across multiple machines.




Figure 3.13: Schematic representation of the flow of data between components of the labscript-suite. Green dashed lines denote threads within a program while blue dashed lines denote a worker process. References to a shot file are passed via ZeroMQ, both between components of the labscript suite and between a component and its worker processes.

BLACS: Hardware Execution Layer

The *Buffered Laboratory Acquisition and Control System* (BLACS) is the interface between the compiled experimental sequence and the physical hardware. For each shot, BLACS loads precompiled instruction buffers and dispatches them to the devices with deterministic timing. The system relies on a multithreaded architecture in which each device - or family of devices - is assigned to an independent control *worker*.

This ensures robust operation and prevents a single hardware failure from blocking the entire experiment. BLACS implements different backend strategies depending on device type:

- **NI-DAQ / National Instruments boards:** Analog and digital channels are programmed via precompiled timing tables with a global hardware clock. This guarantees sub-microsecond jitter and deterministic timing for AOM amplitude ramps, coil currents, shutter control, and trigger signals.
- **DDS and RF sources:** Devices such as the Novatech 409B or the AD9910-based boards are operated through Labscript-provided APIs that write frequency, amplitude, and phase updates into hardware sequences. For fast operations (e.g., frequency chirps for AOD steering), Labscript relies on the hardware's internal ramp engines or FPGA sequencers.
- **Lasers and slow actuators:** Piezo elements, offset locks, temperature controls, or TTL-based laser switches are represented as slow channels. BLACS handles these devices in software-controlled loops that do not require hardware-level determinism.
- **Scientific cameras:** High-speed CMOS and CCD cameras are triggered via TTL pulses generated by NI-DAQ or FPGA boards. BLACS stores metadata for each exposure and passes the resulting data to the analysis stage.
- **Custom hardware and power supplies:** Devices communicating through serial (RS232/USB), Ethernet, or custom protocols are wrapped using Labscript device classes. This approach was used to integrate our magnetic-field power supplies, motorized stages, and the AOD RF drivers.

A key advantage of BLACS is its dual operation mode, combining a fully interactive *manual mode* with the execution *buffered mode* used for running experimental shots. Here, a *shot* refers to a single cycle of an experimental sequence, such as one measurement performed for a specific set of parameters within a larger acquisition campaign.

Between successive shots, the user can directly control any device through BLACS's interactive panels, allowing alignment, debugging, or laser tuning without disrupting the overall workflow. However, in experiments that rely on acquiring large datasets driven sequences -

quite usual for AMO labs - the transition time between consecutive shots can substantially increase the total sequence duration. To address this limitation, we introduce in here a dedicated strategy designed to bypass unnecessary transitions into manual mode, thereby reducing overhead and accelerating high-throughput data collection.

Runmanager: Parameter Definition and Scan Engine

The *Runmanager* GUI provides the parameter layer of the Labscript architecture. Users define numerical, logical, or string parameters that will be substituted into the *python* control script during compilation. This mechanism enables:

- multidimensional parameter scans (frequency sweeps, MOT detuning curves, Rydberg pulses, imaging exposure times);
- automatic generation of sequences for calibration and optimization;
- explicit linkage between metadata and acquired data.

Each parameter is automatically stored in the .h5 (HDF5 format) output file generated for every shot. This guarantees full reproducibility: every dataset is accompanied by the exact configuration used to produce it.

In our implementation, Runmanager controls:

- timing sequences for red/blue MOTs, cooling, imaging, and Rydberg pulses;
- MOT detuning, intensities, and magnetic gradients;
- AWG presets determining AOD dynamics;
- scan variables such as loading time, tweezer power, and repumper settings.
- storage of all auxiliary *python* support codes that contribute to the shot logic, ensuring that the full suite environment is archived together with the experimental parameters.

Through its batch execution system, Runmanager orchestrates entire experimental campaigns without user intervention.

Since usual codes I implemented or fully wrote exceed thousands of lines I will report in this manuscript only the relevant integration points and the overall scheme, omitting lower-level details for clarity.

In List. 1 below we report our integration to generate_code process in `labscript.py` that saves automatically all our codes involved in a shot encoding in its `.h5` dataset when executed by *Runmanager*.

Listing 1: Auxiliary code storage integration in `labscript.py`

```
def save_labscripts(hdf5_file):
    """Writes the script files for the compiled shot to the shot
       file. If 'save_hg_info' labconfig parameter is 'True', will
       attempt to save hg version info as an attribute.
    Args:
        hdf5_file (:obj:'h5py:h5py.File'): Handle to file to save to.
    """
    if compiler.labscript_file is not None:
        script_text = open(compiler.labscript_file).read()
    else:
        script_text = ''
    script = hdf5_file.create_dataset('script', data=script_text)
    script.attrs['name'] =
        os.path.basename(compiler.labscript_file).encode() if
        compiler.labscript_file is not None else ''
    script.attrs['path'] =
        os.path.dirname(compiler.labscript_file).encode() if
        compiler.labscript_file is not None else sys.path[0]
    try:
        import labscriptlib
    except ImportError:
        return
    prefix = os.path.dirname(labscriptlib.__file__)
    for module in sys.modules.values():
        if getattr(module, '__file__', None) is not None:
            path = os.path.abspath(module.__file__)
            if path.startswith(prefix) and (path.endswith('.pyc') or
            path.endswith('.py')):
                path = path.replace('.pyc', '.py')
                save_path = 'labscriptlib/' + path.replace(prefix,
                '').replace('\\', '/').replace('///', '/')
                if save_path in hdf5_file:
                    continue
                hdf5_file.create_dataset(save_path,
                    data=open(path).read())
            with _vcs_cache_lock:
                if not path in _vcs_cache:
```

```

        # Add file to watch list and create its entry
        _file_watcher.add_file(path)
        _file_watcher_callback(path, None, None)
    # Save the cached vcs output to the file.
    for command, info, err in _vcs_cache[path]:
        attribute_str = command[0] + ' ' + command[1]
        hdf5_file[save_path].attrs[attribute_str] = (info
            + '\n' + err)
#####
integrated in :
#####
def generate_code():
    """Compiles a shot and saves it to the shot file.
    """
    [...]
    with h5py.File(compiler.hdf5_filename, 'a') as hdf5_file:
        [...]
        generate_connection_table(hdf5_file)
        write_device_properties(hdf5_file)
        generate_wait_table(hdf5_file)
        save_labscripts(hdf5_file)
        [...]

```

Visualmanager: Timing Visualization and Debugging

Although Labscript code is written as text-based code, *Visualmanager* converts the compiled shot into a graphical timing diagram, making it straightforward to verify the timing logic.

Visualmanager displays:

- all digital transitions (shutters, triggers, AOM switches);
- analog ramps (AOM amplitude, magnetic fields, beam power);
- RF updates for DDS-based devices;
- camera triggers and exposure windows;
- relative timing of all events down to microsecond resolution.

This component is especially valuable when integrating new hardware, ensuring that all trigger lines, ramps, and time-critical sequences behave as intended. In our setup, it was extensively used during the integration of the imaging sequence.

Lyse: Automated Analysis and Real-Time Feedback

The final stage of the Labscript Suite is *Lyse*, the automated analysis environment. After each acquisition, *Lyse*:

1. loads the newly executed .h5 file sent by BLACS;
2. executes one or more user-defined *python* analysis scripts;
3. displays the resulting processed data in real time (e.g. MOT images, atom number, temperature, trap frequencies, fluorescence counts, Rydberg excitation rates, etc.);
4. optionally writes updated experimental parameters for the next iteration, enabling closed-loop optimization.

This allows a broad class of feedback-driven procedures, such as:

- optimization of MOT loading and cooling efficiency,
- adaptive stabilization of tweezer power or beam pointing,
- automated frequency-locking routines or drift compensation,
- feedback-based tuning of Rydberg pulses and excitation conditions.

Lyse supports two complementary modes of operation:

- *Single-shot analysis*, in which each shot is processed independently as soon as its .h5 file is loaded into the *Lyse* pool. This is ideal for real-time monitoring and for feedback loops that require an immediate response.
- *Multi-shot analysis*, in which the same analysis script receives a *list* of shots and can perform correlations, averages, fits over entire scans, or global optimization routines. This mode is crucial for evaluating quantities that cannot be extracted from a single measurement, such as loading statistics, temperature over a parameter sweep, or Rabi oscillations collected across many shots.

In our implementation, we developed a custom multi-shot analysis framework built on a carefully layered try/except architecture. Experimental scans often involve hundreds of sequential shots, and occasional timing irregularities (e.g. delayed file writing, missing datasets, or incomplete metadata) can interrupt the default *Lyse* execution flow.

Our workflow instead:

- attempts to load each shot individually using a protected `try`-block,
- safely skips corrupted or incomplete files while logging the issue,
- preserves the consistency of the analysis arrays (e.g. parameter vectors and data vectors remain aligned),
- proceeds with the remaining valid shots, ensuring that long scans are never interrupted by a single error.

This greatly enhances reliability during overnight or long-duration data acquisition, where intermittent timing glitches would otherwise compromise the entire scan.

Finally, in our implementation each `.h5` file includes not only the numerical analysis results but also the complete *python* analysis scripts executed on itself by Lyse (List. 2). This ensures that both the data and the full processing pipeline are preserved, enabling complete reproducibility and long-term traceability of the experimental workflow.

Listing 2: Analysis code storage integration

```
def saving_script(path):

    with Run(path).open('r+') as shot:
        data_frame = data()
        file_path = os.path.realpath(__file__)
        prefix = os.path.dirname(analysislib.__file__)
        save_path = 'analysislib/' + file_path.replace(prefix,
            '').replace('\\', '/').replace('///', '/')

        with h5py.File(path, 'r') as hdf5_file:
            if save_path not in hdf5_file:
                hdf5_file.create_dataset(save_path,
                    data=open(file_path).read())

    _vcs_cache_rlock = threading.RLock()
```

3.3.2 *Devices Implementation*

Integrating a variety of heterogeneous devices into Labscript requires consistent abstraction layers and dedicated device classes. Below we briefly summarize the strategies we employed.

1. **ABSTRACT DEVICE CLASSES IN PYTHON.** Every routine device - cameras, aom driver, AWG, etc. - is associated with a custom *python* class that encapsulates:

- channel definitions (digital/analog),
- valid operations (ramps, pulses, frequency jumps),
- hardware constraints (update rates, voltage ranges, trigger logic).

2. **CENTRALIZED TIMING MODEL.** All hardware events are scheduled on a single monotonically increasing timeline, ensuring deterministic synchronization. Even devices with their own clocks (DDS or cameras) are synchronized to this timeline through trigger signals.

3. **HARDWARE SEGMENTATION.** Time-critical devices (AOMs, DDS, shutters) use buffer-based execution through BLACS workers, while slow devices (temperature controllers, piezo drivers) are updated asynchronously via software threads.

4. **USE OF INTERMEDIARY DIGITAL TRIGGERS.** Cameras, high-power RF drivers, and coils power-supply are all connected to FPGA-generated TTLs. This makes the experiment immune to possible software jitter.

5. **UNIFIED METADATA AND CONFIGURATION STORAGE.** All parameters passed through Runmanager are written to the same .h5 shot file, ensuring that: every image or analog trace is associated with its configuration; scans are automatically reconstructed; processing pipelines remain reproducible.

6. **DEVICE-AGNOSTIC CODING STYLE.** The Labscript *python* code references logical channels rather than hardware names. This makes it easy to replace devices (e.g. swapping a DDS board or adding a new AOM) without changing high-level code.

Altogether, the Labscript Suite provides a flexible and scalable infrastructure that meets the requirements of our experiment: deterministic timing, robust multi-device integration, an accessible parameter workflow, and real-time data-driven feedback.

3.3.3 *FPGA-Based Control Hardware*

Experiments involving tightly timed sequences - such as Rydberg excitation, or synchronised camera triggers - may require sub-microsecond timing precision. While software-controlled devices (via USB or Ethernet) are sufficient for slow operations, high-bandwidth operations demand deterministic hardware-level control. For this reason, we integrate an FPGA-based system into the Labscript control stack.

System-on-a-Chip FPGA Architecture

Our implementation leverages the architecture presented in Ref. [116], which introduces a hybrid *System-on-a-Chip* (SoC) platform combining:

- a real-time FPGA module responsible for deterministic digital and analog outputs;
- an embedded CPU running Linux, used for high-level coordination, TCP/IP communication, and dynamic sequence loading;
- shared external memory through which the CPU can upload data and access FPGA registers, and the FPGA hardware efficiently loads timing and waveform data without involving the CPU.

This architecture offers several advantages, which are particularly relevant for AMO physics:

- deterministic timing with jitter < 10 ns (FPGA domain);
- flexible reprogramming of experiments without bitstream reload;
- high update rates for DDS, AOMs, and analog outputs;
- high update rates for digital and analog outputs;
- the ability to load multi-shot sequences into onboard RAM for extremely fast repetition rates.

In this system, the FPGA acts as a *shot sequencer*: after receiving compiled instructions from the control computer, it autonomously executes the full timing sequence without further software intervention. This eliminates network latency and operating-system jitter, enabling repetition cycles far faster than ordinary PC-driven control could achieve. In principle, some extremely fast and expensive computers can drive PCIs outputs faster than FPGA, which at the moment can go only at 25-50 MHz speed. Nevertheless, compared to those difficult solutions, the FPGA-SoC is a cheap and much more reliable alternative.

Integration into the Labscript Framework

The FPGA device, written in collaboration with the autor of Ref.[116], implemented in our experiment follows this SoC philosophy (boarded on a Cora Z7 Zynq-7000) and is integrated into the Labscript Suite through a custom *python* device class `FPGA_board` (see GitHub repository²). The device conforms to Labscript's `PseudoclockDevice` API and provides:

- a *python* interface to access FPGA registers and to upload RAM buffers with digital and analog output channel instructions;
- compilation of Labscript instructions into compact FPGA timing tables;
- TCP communication with the embedded CPU, which programs the FPGA;
- automatic inclusion of all FPGA configuration data into the shot's .h5 file.

The uploaded implementation (`FPGA_device.py`) structures FPGA control into three layers:

1. **High-level device model (*python* / *Labscript*)**
Defines abstract operations such as “set DDS frequency”, “play arbitrary waveform”, or “update digital line at time t ”.
2. **Command compilation layer**
Converts Labscript timing instructions into low-level blocks stored in FPGA RAM: (time tag, channel ID, value).
These are sorted for increasing time before transmission.

² <https://github.com/INO-quantum/FPGA-SoC-experiment-control>

3. FPGA hardware engine

Executes the compiled sequence in real time with FPGA-level determinism.

Algorithmic Flow

The execution pipeline for FPGA-based devices in the Labscript Suite follows the sequence below (Alg. 3). This description is device-agnostic and reflects the behaviour implemented in our `FPGA_device.py` driver.

Algorithm 3: FPGA Shot Execution Pipeline

- 1 **Experiment parameters are defined** in Runmanager
 - 2 Labscript **compiles the shot file** into a timing table (`.h5`): the `FPGA_board` class **parses the table** and generates: digital edge sequences, DDS/AWG update lists, and timing words and state-machine entries. These instruction blocks are **packed into binary packets**
 - 3 BLACS **uploads the packets** to the SoC (CPU) via TCP
 - 4 The embedded CPU **configures the FPGA**: writes registers, loads block RAM, and arms internal state machines
 - 5 BLACS issues a **start** command
 - 6 FPGA **executes the full sequence deterministically** in hardware, independent of the host computer
 - 7 Upon completion, FPGA raises a **shot-complete flag** (read by BLACS for synchronization)
-

This architecture allows microsecond timing precision and extremely fast shot-to-shot repetition, since sequence execution is performed entirely in hardware without software latency.

Multi-Channel Control: AOMs, DDSs, Arbitrary Waveforms

The FPGA firmware supports:

- high-speed digital outputs for AOM switching, coils switching and camera triggers;
- DDS frequency and amplitude updates with microsecond latency;
- arbitrary analog waveforms (for trap ramps, magnetic-field shaping, etc.).

The *python* class transparently exposes these as Labscript “channels”:

```
A00=AnalogChannels(name='A00', parent_device=main_board, ...
    rack=0, max_channels = 4)
AnalogOut(name='BigCoilsI', parent_device=A00, connection='0x08')

D02=DigitalChannels(name='D02', parent_device=main_board, ...
    connection='0x05', rack=0, max_channels = 16)
DigitalOut(name='Red_Switch', parent_device=D02, connection=str(0))
```

Multi-Shot Sequences and Safety

A key benefit of SoC-style hardware is the ability to store a *multi-shot run* on the FPGA, allowing the experiment to be repeated hundreds of times with minimal dead time. In our implementation:

- Labscript compiles N shots into N FPGA tables;
- FPGA executes them sequentially without returning to software;
- BLACS monitors flags for errors or end-of-sequence states.

To ensure robustness, all multi-shot analysis scripts in Lyse employ try/except structures to tolerate occasional failed shots (timing issues, camera readout problems, or FPGA timeouts). Failed shots are logged but do not interrupt multi-shot execution.

Advantages of FPGA Integration

The FPGA-based control system provides:

- deterministic and reproducible timing;
- extremely fast shot repetition rates;
- reduced overhead in AOM-scan and tweezer-array experiments;
- high-bandwidth analog and digital updates;
- seamless integration into Labscript with minimal user intervention.

This makes the FPGA a natural backbone for advanced experiments involving dynamic tweezer arrays, Rydberg control pulses, and rapid MOT–tweezer cycles, all of which rely on microsecond-level synchronization.

3.3.4 MOGLabs QRF Driver

The MOGLabs QRF is a four-channel RF synthesizer widely used in AMO experiments for driving AOMs, EOMs, and fast RF control hardware. Starting from the *python*-based control interface provided by the manufacturer, we developed a custom driver that integrates the QRF (–241) module seamlessly into the *Labscript Suite*. The device now exposes its full functionality—static RF generation, table-based frequency and amplitude sequencing, fast power switching, and internal PID stabilization for either frequency or power lock — directly within the BLACS–Runmanager–Lyse workflow. This allows the QRF to be programmed, synchronized, and interrogated using the same abstraction layers as the rest of the experimental hardware, enabling deterministic timing, automated parameter scans, and real-time analysis feedback. The implementation is organized into three components:

1. Labscript device classes, which define the logical interface and timing model;
2. BLACS GUI components, which provide real-time interactive control;
3. a BLACS worker, responsible for low-level TCP communication with the hardware.

Labscript Device Layer

Each RF channel is represented by a QRF_DDS object, supporting two operating modes:

- *Static mode*: a fixed RF tone is produced throughout the experimental sequence, with switching performed by fast TTL gates for nanosecond timing accuracy;
- *Table mode*: a time-ordered sequence of RF instructions is programmed into the QRF’s internal memory, enabling deterministic multi-point frequency, amplitude, or phase modulation synchronized with the pseudoclock.

The main device class, MOGLabs_QRF, aggregates these channels and implements the interface with Labscript’s pseudoclock system.

During compilation, `generate_code()` writes either:

- `STATIC_DATA` groups (final RF values), or
- `TABLE_DATA` groups (full dynamic RF tables)

into the shot's `.h5` file. Timing constraints for table-triggered operation are automatically communicated to the compiler.

BLACS GUI Layer

A dedicated BLACS tab is generated, providing the user with interactive control of:

- frequency, amplitude, and phase for each channel (real-time updates in manual mode);
- power state via a custom “*Signal / Amplifier / Both*” switching widget;
- PID control parameters: proportional, integral, derivative gains, pregain, polarity, and real-time error readout.

The GUI implements enhanced remote-value checking to ensure that the displayed state matches the hardware state. This includes robust reconnection logic, recovery from incomplete table executions, and safe fallback to static mode.

BLACS Worker: Low-Level Communication

The worker class `MOGLabs_QRF_Worker` handles all communication with the hardware using the MOGLabs ASCII protocol. Its main responsibilities include:

- detection and initialization of channels on startup;
- recovery from communication errors and missing devices;
- configuration of RF output during `transition_to_buffered()`;
- application of all manual controls in `program_manual()`;
- safe return to static mode in `transition_to_manual()`;
- reinsertion of PID states and setpoints if active before the shot.

During shot execution, the worker:

1. reads the .h5 file generated by Labscript;
2. loads table entries (for table mode) or applies final static values;
3. arms the device for hardware triggering;
4. ensures correct timing resolution based on operating mode;
5. enables or disables PIDs as required.

A “smart programming” cache avoids redundant communication when consecutive shots reuse identical RF settings.

Key Features

The driver provides:

- full compatibility with Labscript’s event-based timing model;
- nanosecond-precision static gating and microsecond-resolution table sequencing;
- integrated PID amplitude stabilization with real-time tuning;
- robust communication and automatic recovery from transient errors;
- synchronized multi-channel RF control, including frequency ramps and dynamic AOM drive.

This implementation therefore makes the MOGLabs QRF a fully programmable device within the *Labscript Suite*, supporting both interactive alignment tasks and high-repetition automated sequences essential for cold-atom experiments.

3.3.5 Basler ACE Pylon Camera

High-performance imaging is a fundamental component of cold-atom experiments, particularly for absorption and fluorescence detection. In our setup, we integrated a Basler ACE camera using a custom Labscript device built on top of the Pylon API ³.

³ <https://www.baslerweb.com/en/software/pylon/>

The implementation consists of two coordinated modules, following the standard Labscript Suite architecture:

1. a `labscript_devices` class, defining the experiment-time interface and device properties;
2. a BLACS Worker providing real-time communication with the camera hardware.

In `labscript_devices.py`, the camera is registered as a subclass of the generic Camera device:

```
class MyPylonCamera(Camera):  
    description = 'MyPylon Camera'
```

This definition exposes the camera to the Labscript compiler, enabling configuration of exposure time, trigger behaviour, region of interest (ROI), and acquisition timing within an experimental shot.

The real hardware logic is implemented in the worker class called `MyPylonCameraWorker`, which inherits from `IMAQdxCameraWorker`. It instantiates a `MyPylon_Camera` object that wraps all Pylon API calls.

Upon initialisation, the worker:

- loads the Pylon API dynamically (to avoid import errors on systems without the library);
- identifies the camera by its serial number through `TlFactory`;
- opens the device, registers software triggering, and caches the GenICam node map.

Camera configuration is performed by writing directly to the GenICam parameter tree. Because ROI settings require a specific order (width before offset or vice-versa, depending on boundaries), the worker automatically resolves the correct update sequence:

- ROI-x and ROI-y parameters are updated atomically and in the correct order to avoid GenICam exceptions;
- all other parameters are set through an unified `set_attribute(name, value)` interface;
- attribute discovery is performed using visibility and access-mode filters (Beginner/Expert/Guru and RW/RO).

The camera supports three operational modes:

1. **Single-shot acquisition** using `GrabOne`;
2. **Continuous acquisition** through the Pylon grab engine (`GrabStrategy_LatestImageOnly`);
3. **Buffered acquisition** for multi-image sequences.

Buffered acquisition is implemented through an iterative retrieval loop:

- images are collected with `RetrieveResult`;
- timeouts are handled automatically (retry on timeout);
- acquisitions may be aborted through a shared flag, guaranteeing that BLACS can safely interrupt a camera scan without freezing the worker.

At the end of each shot:

- the worker stops the grab engine,
- flushes the internal image queue,
- and finally closes the camera connection when BLACS is shut down.

This ensures that long sequences of shots proceed reliably without producing dangling buffers or stale handles.

Within the Labscript Suite:

- **Runmanager** sets camera parameters (exposure, gain, ROI, and frame counts);
- **BLACS** triggers acquisitions at precise times within the experimental sequence;
- **Lyse** receives the resulting datasets from the `.h5` shot file.

The Pylon-based implementation provides low-latency triggering, deterministic frame acquisition, and stable repetition-rate performance—critical for experiments involving:

- high-resolution absorption imaging,
- fluorescence detection from single atoms,
- multi-shot imaging sequences for temperature or trap-frequency measurements.

SUMMARY. The custom Pylon camera driver is integrated into the Labscript ecosystem, providing:

- robust ROI and parameter handling through the GenICam interface;
- deterministic acquisition both in continuous and triggered modes;
- safe multi-shot buffered acquisition with abort handling;
- complete compatibility with Labscript timing, Runmanager parameter injection, and Lyse analysis pipelines.

This implementation allows our imaging system to operate reliably at high shot rates and provides the flexibility required for advanced experimental protocols.

3.3.6 Hamamatsu ORCA Quest Camera

The Hamamatsu ORCA Quest C15550-20UP camera is integrated into the experimental control system through a custom labscript device built on top of the DCAM-SDK (accessed via the *python* wrapper used in our laboratory). The implementation consists of two tightly coupled layers:

1. a low-level hardware interface (DCAM_Camera), responsible for all DCAM API calls, buffer handling, and frame acquisition;
2. a BLACS worker class (DCAMCameraWorker), which manages timing, synchronization, and the transfer of acquired images into the .h5 shot file.

Beyond its software integration, the ORCA Quest camera plays a fundamental role in the experiment itself. In particular, its high quantum efficiency, low readout noise, and fast frame-transfer capabilities are essential for resolving single-atom fluorescence with high fidelity. The ORCA Quest camera features a peak quantum efficiency exceeding 70% (and up to $\sim 90\%$ in the visible range), combined with a readout noise at the level of $\sim 0.9e^- \text{ RMS}$ in low-noise modes.

Together with frame-transfer operation and typical full-frame acquisition rates of 50–100 frames per second (and significantly higher for cropped regions of interest), this performance enables reliable discrimination between occupied and empty tweezers with exposure times of only a few milliseconds. Such short imaging times are crucial both for minimizing atom loss and heating during fluorescence detection, and

for operating the experiment at high repetition rates compatible with large-scale data acquisition and real-time feedback.

Single-atom fluorescence detection requires high quantum efficiency at the 461 nm imaging wavelengths, low readout noise, and fast, reliable frame transfer. The ORCA Quest platform combines deep-cooled low-noise CMOS architecture with high sensitivity, enabling the clear discrimination between photon-shot-noise-limited single-atom signals and background scatter. Its rapid readout and subsystem-synchronized triggering also allow repeated imaging cycles within a single experimental shot, essential for cooling diagnostics, trap loading characterization, and atom-by-atom rearrangement procedures.

Together, the hardware capabilities of the ORCA Quest camera and the custom *Labscript* integration ensure high-fidelity, low-latency fluorescence imaging, which is crucial for resolving individual atoms in optical tweezers and for enabling the feedback algorithms used to prepare defect-free arrays.

DCAM Backend: Low-Level Hardware Control

The `DCAM_Camera` class provides a complete *python* interface to the Hamamatsu DCAM API. Upon initialization, the class:

- initializes the DCAM library and enumerates all detected cameras;
- selects the camera by serial number, ensuring deterministic device assignment;
- retrieves the full list of readable and writable camera properties and stores them in a *python* dictionary for fast access;
- applies a default configuration (exposure, trigger mode, readout format, etc.) prior to BLACS use.

The class exposes standard camera operations required by *Labscript*:

- **`snap()`**: acquire a single frame using internal triggering;
- **`configure_acquisition()`**: allocate buffers, choose continuous or single-frame mode, and start capturing;
- **`grab()`** and **`grab_multiple()`**: retrieve buffered frames by index or as a sequence, with timeout and error handling;

- **stop_acquisition()** and **abort_acquisition()**: cleanly shut down acquisition and release buffers.

Buffered acquisition is implemented through a loop that uses fast, repeated DCAM timeout calls. Instead of blocking, the code continuously polls both the DCAM API and an internal `_abort_acquisition` flag, allowing BLACS to interrupt long or faulty acquisitions gracefully.

BLACS Worker: Threaded Acquisition and HDF5 Integration

The `DCAMCameraWorker` handles all coordination with BLACS and the shot compilation pipeline. During the *transition-to-buffered* stage, the worker:

1. reads the list of exposures (names, frametypes, timestamps) from the shot's `.h5` file;
2. configures the camera with exactly the number of required buffers;
3. launches a dedicated acquisition thread that repeatedly calls `grab()` or `grab_multiple()` from the DCAM backend.

A key feature of our implementation is that each frame is written into the HDF5 file *immediately upon acquisition*. This is handled by the method `save_image_now()`, which:

- creates the appropriate HDF5 group and dataset for each exposure;
- writes the image with gzip compression;
- attaches image metadata needed by downstream analysis (image class, version, white-is-zero flag, etc.).

Saving frames on-the-fly provides two major advantages. First, it makes the acquisition pipeline resilient to abnormal camera behavior or partial-image delivery: even if a shot fails midway, all successfully recorded frames are already written to disk. Second, it enables real-time analysis feedback, dynamic platform reordering, and error-correction strategies, since the analysis routines can begin processing data as soon as each frame becomes available.

Transition to Manual Mode and GUI Display

When BLACS returns to manual mode, the worker:

1. joins the acquisition thread with a timeout and raises a controlled error if frames were not acquired as expected;
2. stops the camera and releases DCAM buffers cleanly;
3. forwards stacked images to the BLACS GUI for live display (when possible);
4. restores the manually-defined camera attributes.

This ensures that the camera can seamlessly switch between single-shot experimental execution and interactive manual use, preserving DCAM state and preventing buffer conflicts.

Robustness Features

Several layers of protection are implemented:

- **Timeout-controlled acquisition** prevents camera hangs from blocking BLACS.
- **Abort flags** allow BLACS to interrupt multi-image acquisitions immediately.
- **Automatic HDF5 writing** guarantees that partial datasets are never lost.
- **Safe reconfiguration logic** only reprograms attributes that differ from the cached configuration, reducing latency and avoiding unnecessary API calls.

Overall, this integration provides a high-speed, reliable, and fully synchronized camera control interface, compatible with both single-number extractions (e.g., MOT population) and high-throughput multi-image sequences such as fluorescence scans, Rydberg excitation detection, and tweezer-array diagnostics.

3.3.7 *Spectrum AWG Card*

A key component of our control infrastructure is the high-performance Arbitrary-Waveform Generator (AWG) based on a Spectrum M4i series card. This device serves as the primary driver of the AODs used to generate and steer the optical tweezers. The AWG enables deterministic, low-latency control of frequency and amplitude waveforms, allowing smooth and reproducible atom transport while maintaining high experimental repetition rates.

Its tight integration with the control software is essential for the atom-reordering phase, where a closed feedback loop (see Sec. 3.3.8) interconnects image acquisition, real-time image analysis, reordering algorithms, and the precise sequencing of tweezer trajectories. To integrate the AWG into the *Labscript Suite*, we developed two tightly coordinated modules:

1. a dedicated BLACS worker (`SpectrumAWGWorker`), which exposes the card to the experimental-control layer;
2. a low-level hardware abstraction layer (`SpectrumCard`), which provides a complete and pythonic interface to the vendor's C API.

Together, these modules provide deterministic, high-bandwidth waveform delivery, fast hardware triggering, and support for advanced sequence-replay modes, all of which are essential for dynamical tweezer experiments and rapid reordering protocols.

Worker Architecture

The `SpectrumAWGWorker` class is responsible for initializing the device, managing channel configuration, writing samples to memory, and creating complex sequence-replay instructions. Upon initialization, the worker:

- opens a handle to the AWG via the `SpectrumCard` class;
- configures internal or external clocking (with MHz–GHz rates);
- enables the subset of channels requested from BLACS;
- sets output amplitude, filtering options, and generation modes;
- configures external triggering thresholds and behaviour;
- divides the on-board memory into user-selected segments for sequence replay.

Each AWG channel operates independently and is controlled by dedicated thread-safe methods. Manual (single-tone) and buffered (sequence-replay) modes are supported. For manual generation, the worker synthesizes custom sine tones in software using helper functions such as:

- `generate_single_tone()`: basic sine generation at fixed amplitude;
- `generate_manual_tone()`: two-channel amplitude-controlled tones;
- `generate_multi_tone()`: normalized multi-sine combinations.

SMART MEMORY CACHE To minimize transfer time, we implemented a “smart programming” layer: each instruction is hashed, and identical segments are written to the AWG memory only once. When a repeated instruction is encountered in a future shot, the worker simply references the existing memory segment. This reduces DMA transfers by orders of magnitude during multi-shot runs.

REORDERING-MODE THREAD For fast atom-rearrangement experiments (see Sec. 3.2.5), the worker provides a dedicated background thread that parses reordering instructions from the .h5 file and streams only the required waveform segments to the card.

Because the array geometry may change between shots, the worker continuously monitors the `geometry_path` parameter defined in the `globals` Runmanager settings. This parameter points to the file describing the current tweezer geometry and the corresponding set of allowed reordering paths. The internal cache is reset only when a change in the geometry is detected, thereby avoiding unnecessary recomputation while ensuring consistency with the active experimental configuration.

Transition to Buffered Mode

During experiment compilation, BLACS calls `transition_to_buffered()`, which:

1. stops the card if running (e.g. from manual mode);
2. sets sample rates, trigger configuration and replay mode;
3. parses all waveform samples stored in .h5;
4. uses the smart cache to determine whether each waveform must be transferred to memory or can be reused;

5. programs the sequence steps using the card's native sequence registers;
6. starts the card and enables triggering.

Sequence steps are encoded in 64-bit registers containing:

- memory segment index,
- next-step pointer,
- loop count,
- optional “end-on-trigger” flag,
- end-of-sequence flag.

This mechanism allows complex multi-step waveform choreography entirely on-board, with deterministic sub-nanosecond timing.

Driver Layer: The SpectrumCard Class

The SpectrumCard module provides a complete and clean interface to the Spectrum M4i driver API. It abstracts all hardware registers into *python* functions, including:

- card discovery, serial number, PCI capability queries;
- configuration of internal / external clocks;
- control of channel enable masks, amplitudes, filters and output modes;
- configuration of replay modes (single, multi, gated, FIFO, DDS);
- full access to trigger masks, levels, AND/OR logic, and rearming modes;
- writing waveform segments to memory using DMA transfers;
- configuring sequence steps and loop conditions;
- card start, stop, wait-ready, force-trigger, and reset commands.

The driver ensures that all numerical parameters are validated, handles conversion to appropriate C types, and provides human-readable information for debugging.

Waveform Generation Workflow

Algorithm 4 summarizes the full workflow for generating and executing a waveform sequence.

Algorithm 4: AWG Sequence Programming Workflow

```

1 Initialize card (clock, channels, trigger, memory);
2 for each sequence step s in shot file do
3   | extract waveform label → compute hash;
4   | if hash in smart cache then
5   |   | reuse existing memory segment
6   | else
7   |   | synthesize / load waveform samples;
8   |   | DMA transfer to AWG memory;
9   |   | update smart cache;
10  | program sequence-step register (segment, next, loops)
11 Start card and enable trigger

```

Table-Based Signal Generation for AOD Tweezer Control

Dynamic control of optical tweezers generated with AODs requires the synthesis of precisely shaped RF waveforms, whose instantaneous frequency determines the deflection angle of the tweezer beam, while the amplitude determines the trap depth. To implement arbitrary 2D tweezer motion - including point-to-point atomic transport, line translations, oblique motion, and composite standing-wave patterns - we developed a flexible AWG driver that constructs waveforms directly from *tabled motion schedules*. These schedules are stored in external .csv files and specify, as a function of time, the desired displacement and trap amplitude. We are currently testing the theoretical work of Refs.[117].

This design achieves several goals:

- it decouples trajectory design from hardware implementation;
- it ensures that all AWG waveforms obey the device's strict memory and sampling-alignment constraints (4096-bit sample granularity);
- it enables full reproducibility: each motion sequence can be regenerated from its table.

GEOMETRY DEFINITION. The set of tweezer coordinates is loaded either from a uniform square grid (for regular arrays) or from a geometry file (for arbitrary graph layouts). The coordinate dictionary

$$\text{points_dict} = \{P_i \mapsto (x_i, y_i)\}$$

provides the mapping between point index and physical position (in μm). Edges connecting sites, used for reordering (Sec. 3.3.8), are likewise loaded for graph-based motion planning.

FREQUENCY-POSITION CALIBRATION. AOD calibration gives linear conversion between position and driving RF:

$$f_x(x) = f_{0,x} + \alpha_x x, \quad f_y(y) = f_{0,y} + \alpha_y y,$$

where f_0 are the center frequencies and α the scanning coefficients ($\text{Hz}/\mu\text{m}$). Horizontal, vertical, and oblique moves are automatically detected from the start and end coordinates, and the corresponding mapping function is selected.

AMPLITUDE CORRECTION: LUT AND GAUSSIAN COMPENSATION. Because the AOD diffraction efficiency, beam shape, and overall optical throughput vary across the scanning range, the RF amplitude must be corrected in order to maintain a uniform optical trap depth. In our implementation, we apply two complementary correction mechanisms that address distinct physical effects:

1. **Gaussian (device-level) correction**, which models the intrinsic frequency response of the AOD. This correction captures the smooth roll-off in diffraction efficiency that arises from the finite acoustic bandwidth and electromechanical response of the AOD itself. It is well described by a Gaussian-like dependence on the RF frequency,

$$C_{\text{Gauss}}(f) = a + (1 - a) \left[1 - \exp \left(-\frac{1}{2} \left(\frac{f - f_0}{\alpha \sigma} \right)^2 \right) \right], \quad (3.15)$$

where σ is calibrated independently for each scanning axis and reflects the intrinsic AOD performance rather than optical effects.

2. **LUT-based (optical-path) correction**, which compensates for residual spatial variations originating from the optical beam itself and the downstream optical system. This correction is obtained from

a measured intensity map acquired by scanning a single optical tweezer over the full accessible x - y plane. The resulting map is cropped and normalized to generate a spatially varying amplitude correction (LookUp Table - LUT), accounting for aberrations, beam clipping, and nonuniform transmission along the optical path.

The total correction applied to the RF envelope is given by

$$C(f) = C_{\text{Gauss}}(f), C_{\text{LUT}}(f). \quad (3.16)$$

While these two corrections partially overlap in their effect, they originate from different physical mechanisms: the Gaussian term captures the intrinsic AOD frequency response, whereas the LUT correction addresses beam- and optics-related inhomogeneities. Using both therefore represents a controlled and slightly overcomplete approach that ensures robust uniformity of trap depths across the full scanning range, which is crucial for high-fidelity atom transport.

CSV-BASED MOTION SCHEDULES. Each motion sequence is defined starting from a table:

$$\{ t_i, x_i, A_i \}_{i=1}^{N_{\text{steps}}},$$

where:

- t_i is the timestamp of the i -th interpolation knot,
- x_i is the displacement (converted into position),
- A_i is the trap amplitude (converted into RF voltage).

Different CSV files (e.g. T10.csv, T20.csv, ...) correspond to the theoretical optimal motion profiles designed for different transport durations.

SAMPLE GENERATION. For a given move, the AWG driver performs the following steps:

1. converts start and end positions into their corresponding frequencies.
2. interpolates frequency and amplitude linearly between table timestamps.
3. applies Gaussian and LUT amplitude corrections to each instantaneous frequency.

4. integrates frequency to obtain phase:

$$\phi(t) = \int 2\pi f(t) dt.$$

5. generates the RF waveform:

$$s(t) = A(t) \sin(\phi(t)).$$

6. quantizes to 16-bit signed integers.
7. aligns the total length to the AWG's 4096-bit sample memory granularity.
8. upload to the appropriate channel (horizontal or vertical).

Horizontal and vertical waveforms are generated independently, allowing for synchronized oblique trajectories by combining both axes. For purely horizontal (vertical) moves, the orthogonal axis is held at a static standing-wave frequency, again amplitude-corrected using LUT and Gaussian compensation.

STANDING AND COMPOSITE WAVEFORMS. The same framework also supports:

- **multi-frequency standing waves** that address several tweezers simultaneously;
- **equalized standing waves**, produced in the frequency domain to maintain consistent amplitude across spectral components;
- **moving-standing patterns**, where a single or a multiple tweezers moves while others remain fixed.

These features are essential for the implementation of reordering operations (Sec. 3.3.8) and programmable trap configurations.

FULL INTEGRATION WITH THE REORDERING ENGINE. Once the reordering algorithm produces a list of elementary moves $\{i \rightarrow j\}$ (Sec. 3.3.8), each move is mapped onto a tabled schedule through:

$$\text{movingTweezer}(P_i, P_j),$$

which triggers the AWG driver to synthesize the waveform and send it to the hardware.

Composite line shifts (`shiftingTweezerLine`), multi-tweezer standing patterns (`standingTweezer`), and oblique motion (`movingSamples`) are all implemented through the same unified waveform engine.

In summary, this table-driven AWG architecture provides a powerful and flexible toolkit for high-fidelity atom transport, enabling the deterministic rearrangement of arbitrary tweezer geometries required for defect-free quantum registers.

Manual Mode

In manual mode, the worker bypasses the sequence engine and directly streams one or more synthesized tones. This mode is used for:

- trap alignment,
- frequency calibration,
- imaging-beam optimization,
- debugging DDS/AOD behaviour.

Switching between manual and buffered modes occurs without reinitializing the card, ensuring fast experimental turnaround.

Summary

The combined `SpectrumAWGWorker` and `SpectrumCard` framework provides:

- robust, thread-safe integration with Labscript BLACS;
- GHz-sample-rate waveform generation with deterministic timing;
- fast dynamic reprogramming using smart memory caching;
- a complete *python* interface to all low-level hardware features;
- support for reordering and multi-tone tweezer-driving protocols.

This infrastructure underpins the entire optical-tweezer control architecture, from static trap arrays to high-speed Rydberg pulse shaping and dynamic atom rearrangement.

3.3.8 Feedback Strategy

Real-time feedback is essential for operating a neutral-atom quantum simulator, where each experimental realization produces a stochastic loading pattern that must be corrected before coherent many-body dynamics can be executed. Our control system implements a low-latency feedback loop that leverages three key capabilities:

1. on-the-fly frame saving by the fluorescence camera (via the ORCA Quest-implemented `save_image_now` functionality),
2. background AWG worker threads capable of issuing rapid trap-motion updates,
3. a fast connectivity-aware reordering algorithm that computes the atom motion sequence directly from the fluorescence image. [118]

Together, these elements enable deterministic initialization of arbitrary tweezer arrays starting from a probabilistically-loaded distribution.

Real-Time Acquisition and Event Triggering

During imaging, the camera device writes each acquired frame directly to disk using a frame-by-frame `save_image_now` mechanism. This design provides two key advantages:

- *Resilience to partial or failed shots*: even if the experiment aborts prematurely, all successfully acquired images are already safely stored.
- *Immediate availability for processing*: images can be accessed by an independent thread as soon as they appear on disk, without waiting for the shot to finish.

A dedicated monitor thread continuously polls the experiment directory and detects the appearance of new shots or new frames. When the relevant fluorescence frame is written to disk, the thread loads it safely (using retry-aware file access to avoid race conditions) and triggers the reordering stage discussed below.

Reordering Algorithm

Once a fluorescence image is available, atom occupations are extracted by integrating the signal around each known tweezer center, producing the Boolean vector

$$\{n_i\} \in (0,1)^N,$$

which specifies which traps are filled. Because loading is probabilistic, this pattern rarely matches the desired target configuration required for coherent many-body experiments. To correct this, we use the graph-based reordering algorithm implemented in `reordering.py`. The tweezer array is modeled as a graph $G = (V, E)$, where each trap corresponds to a vertex $i \in V$, and allowed single-atom moves define the edges $(i, j) \in E$.




Figure 3.14: Obstacle-based atom reordering protocol. The shaded central square defines the target trap region, whose sites (labeled 1–9) constitute the final node list used by the priority algorithm. In orange all the shortest paths for each free atom to reach trap n. 5. (i) The closest free atom (red-path) is routed toward the target site through the shortest allowed path; when an occupied site is encountered, it is treated as an obstacle. (ii) The obstacle strategy assigns the encountered atom to the target trap, while the original atom is redirected into the obstacle’s initial position, effectively performing an exchange. (iii) Repeating this procedure yields a collision-free sequence of local moves that fills the desired region (shaded) while preserving single-occupancy constraints.

OCCUPATION EXTRACTION. For each site, the fluorescence signal is classified as *occupied* or *empty*. The set of occupied sites is mapped onto a canonical target order, producing a list (called *priority*) target destinations[118]. Traps already occupied after random loading are marked

as completed and removed from the active target list. The remaining empty traps are ordered by priority, starting from the most central sites. Reordering begins with the highest-priority empty trap; in the example of Figure 3.14, trap 5 must be filled. The algorithm selects the closest available atom, indicated by the red path, while the orange paths denote the shortest paths to other, more distant free atoms. During transport, any encountered occupied site is resolved via the obstacle strategy, ensuring a collision-free sequence of local moves that incrementally fills the target region.

PATH PLANNING. To move an atom from site i to site j , the algorithm computes a shortest allowed path $i \rightarrow j$ while avoiding all currently occupied sites except the target. When an occupied site is encountered along this path, the algorithm applies an obstacle strategy: the encountered (obstacle) atom is reassigned to the target trap j , which is scheduled to be filled, while the atom originating from i is moved into the site previously occupied by the obstacle atom. This exchange-based resolution preserves single-occupancy constraints and avoids collisions without requiring temporary storage in auxiliary sites. The procedure results in a consistent, collision-free sequence of elementary one-step translations,

$$\text{move_}i_j,$$

where each move represents the displacement of a single atom between adjacent traps. The procedure is displayed in Figure 3.14.

OUTPUT. The resulting move list is written into the shot's .h5 file under a dedicated dataset (e.g. /reordering_sequence/moves), ensuring complete traceability and reproducibility.

Trap-Motion Execution via AWG Threads

The reordering sequence is executed by the hardware responsible for shifting the optical tweezers, typically an arbitrary waveform generator (AWG). The AWG interface runs inside a dedicated BLACS worker thread, which:

- remains active independently of Labscript compilation,
- receives asynchronous commands from the reordering monitor,
- generates rapid frequency-phase trajectories to translate tweezers between adjacent sites.

When the reordering dataset becomes available, the AWG thread reads the movement instructions, parses the atomic operations, and issues the necessary waveform updates. Because the trap motion is executed at the hardware level, this procedure fits well within the typical inter-shot dead time and does not slow down the experimental cycle.

Closed-Loop Repetition and Deterministic Initialization

The feedback sequence - load, image, detect, reorder - can be repeated until the desired target array is realized. After each reordering step, an additional fluorescence image may be taken to verify the success of the relocations. If residual defects remain, the cycle is repeated using the updated occupancy pattern. This closed-loop strategy converts a randomly-loaded distribution into a deterministic initial state, enabling arbitrary real-space configurations for quantum simulation, such as:

- fully filled registers,
- defect-engineered patterns,
- geometries with structured boundaries or sublattices.

It is this fast, low-latency feedback loop that ultimately enables reliable and reproducible preparation of specific atomic configurations on demand.

3.3.9 *Independent Lab Processes*

Despite its versatility and effectiveness, our control system — the *Lab-script Suite* — does not cover all the automatized tasks that must be run within the apparatus-wide process control infrastructure. Indeed, several independent subsystems are employed to supervise the critical components of the experiment that operate continuously, irrespective of whether scientific shots are being executed. These auxiliary systems constitute what is traditionally referred to as a *process control layer*: they maintain stable operating conditions, ensure hardware protection, and enforce safe behaviour of the apparatus at all times.

Most of these process-control functions are implemented as *hardware interlocks*. Their role is to monitor potentially hazardous or failure-prone subsystems and, upon detecting abnormal conditions, shut down the affected hardware and drive the apparatus into a safe state. While Labscript programs may read the state of the interlocks, they never issue commands to them, ensuring strict one-way communication and preventing software-level failures from propagating into critical safety systems.

In addition to interlocks, several continuous-control processes run independently of Labscript. Examples include laboratory temperature stabilization (implemented both at the room scale and using dedicated optical table enclosures), automated daily oven warm-up scripts executed remotely, laser frequency-locking and power-stabilization and various slow-control routines designed to maintain optimal experimental conditions even when no sequence is running. Together, these systems complement the Labscript Suite: Labscript governs fast, deterministic shot-level control, whereas the parallel process-control layer ensures robustness, stability, and long-term protection of the experimental apparatus.

CHARACTERIZATION OF MACHINE PERFORMANCE

4.1 TWO-STAGE MOTS OF STRONTIUM

This chapter is based on the publication:

V. Giardini, L. Guariento, A. Fantini, S. Storm, M. Inguscio, J. Catani, G. Cappellini, V. Gavryusev, and L. Fallani:

“Single Strontium Atoms in Optical Tweezer Arrays for Quantum Simulation,” MDPI Atoms 2026, 14(1), 1, [doi:10.3390/atoms14010001](https://doi.org/10.3390/atoms14010001).

To cool and trap strontium atoms, we employ a two-stage magneto-optical trap (MOT) (Sec. 2.1.3) sequence that relies on two optical transitions with very different linewidths. The first stage uses the broad $^1S_0 \rightarrow ^1P_1$ transition at 461 nm (blue MOT), which provides a large capture velocity and enables efficient loading from a thermal atomic beam. Its drawback is the relatively high Doppler limit ($T_D \approx 770 \mu\text{K}$). Further cooling is achieved with the narrow intercombination transition $^1S_0 \rightarrow ^3P_1$ at 689 nm (red MOT), which allows us to reach microkelvin temperatures and high atomic densities. However, the extremely small capture velocity of this transition prevents direct loading from the blue MOT.

To bridge this gap, we introduce an intermediate broadband red MOT stage, in which we artificially broaden the laser spectrum. This step ensures efficient transfer from the blue MOT to the single-frequency red MOT. As a result, we obtain cold and dense ensembles in the single-frequency red MOT, with phase-space densities high enough to enable efficient loading into optical tweezers.

Atom number and temperature

A quantitative characterization of a MOT requires determining both the number of trapped atoms and the temperature of the cloud.

These quantities are obtained through absorption imaging of the expanding cloud after a variable time of flight (TOF).

Absorption imaging is a standard diagnostic tool in ultracold-atom experiments, as it directly reveals the spatial distribution of the sample and enables a reliable extraction of its size, atom number, and temperature [52, 119]. The absorption cross section is

$$\sigma(\Delta, s) = \frac{3\lambda^2}{2\pi} \frac{1}{1 + 4\Delta^2/\Gamma^2} \frac{1}{1 + s}, \quad (4.1)$$

where Δ is the detuning, s the saturation parameter, λ the probe wavelength, and Γ the natural linewidth. For resonant ($\Delta = 0$) weak intensity ($s \ll 1$) imaging, it reduces to $\sigma = 3\lambda^2/2\pi$. In the weak intensity regime, the attenuation of the probe beam obeys the Beer–Lambert law,

$$I_t(x, y) = I_0(x, y)e^{-\alpha(x, y)}, \quad \alpha(x, y) = \sigma \tilde{n}(x, y), \quad (4.2)$$

with $\tilde{n}(x, y) = \int n(x, y, z) dz$ the density integrated along the imaging axis z . Hence the column density is

$$\tilde{n}(x, y) = -\frac{1}{\sigma} \log \left(\frac{I_t(x, y)}{I_0(x, y)} \right). \quad (4.3)$$

For the measurements reported here, we use a resonant 461 nm probe beam addressing the $^1S_0 \rightarrow ^1P_1$ transition. After abruptly switching off the MOT beams and the magnetic-field gradient, the atoms are allowed to expand freely and fall under gravity. A short probe pulse is then applied, and the shadow left on the probe beam is imaged onto Ace camera. Each experimental cycle records three images in sequence:

- **Signal (S):** probe beam transmitted through the atomic cloud;
- **Probe (P):** probe beam without atoms, giving the incident intensity profile;
- **Background (B):** camera background acquired without atoms and without probe light.

According to Eq. 4.3, the number of atoms in the cloud can be obtained from [119]:

$$N = -A \sum_{x,y} \frac{1}{\sigma} \log \left(\frac{S_{x,y} - B_{x,y}}{P_{x,y} - B_{x,y}} \right), \quad (4.4)$$

where A is the pixel area, σ is the absorption cross section, and $S_{x,y}$, $P_{x,y}$, $B_{x,y}$ are the pixel intensities of the three images.




Figure 4.1: Absorption image of the blue MOT after $300\ \mu\text{s}$ TOF. **Left:** experimental image with one-dimensional cuts along the x and y axes (blue) and Gaussian fits (red). **Right:** two-dimensional Gaussian fit used to extract atom number and cloud sizes. From the fit we obtain $N = 1.25 \times 10^7$ atoms, a peak density of $3.8 \times 10^9\ \text{cm}^{-3}$, and cloud sizes $\sigma_x = 1.64(2)\ \text{mm}$ and $\sigma_y = 1.69(2)\ \text{mm}$. The cloud profile has been rotated while performing the fitting using Eq. 4.6.

Rather than performing the explicit integration as from Eq. 4.4, we extract the MOT **total atom number** N_{atoms} by fitting the column density with a two-dimensional Gaussian (with σ_x^2 , σ_y^2 variances), which after integration yields

$$N_{\text{atoms}} = 2\pi \sigma_x \sigma_y \frac{n_0}{\sigma}, \quad (4.5)$$

where n_0 is the fitted optical density peak and σ is the atomic resonant absorption cross section as seen in Eq. 4.1. This expression is applicable to both the blue and the red MOT, whose thickness does not exceed the regime in which this approximation remains valid ($\text{OD} < 1$).

In our experiment the absorption images are taken $300\ \mu\text{s}$ after switching off the trapping light and the magnetic gradient. This delay ensures complete suppression of the magnetic field and avoids residual Zeeman shifts that would affect the absorption cross section. An example image of the blue MOT together with the fit is shown in Fig. 4.1. To determine the MOT **temperature** T_{MOT} , we measure the expansion of the cloud for different TOFs, as shown in Fig. 4.2 for the red MOT.

Well above the quantum-degenerate regime, the velocity distribution is Maxwell–Boltzmann, and the cloud width at TOF t follows [119]:

$$\tilde{n}(x, y, t) \propto \exp \left[-\frac{(x - x_0)^2}{2\sigma_x^2(t)} - \frac{(y - y_0)^2}{2\sigma_y^2(t)} \right], \quad (4.6)$$

$$\sigma_i^2(t) = \sigma_{0,i}^2 + \sigma_v^2 t^2, \quad \sigma_v^2 = \frac{k_B T_{\text{MOT}}}{M}, \quad (4.7)$$

where M is the atom mass, k_B is the Boltzmann constant, $\sigma_{0,i}$ is the initial cloud size and σ_v the thermal velocity spread. From the cloud expansion along the two axes, we extract $\sigma_{v,x}^2$ and $\sigma_{v,y}^2$, which agree within uncertainties. We therefore obtain a more accurate temperature from their average,

$$\sigma_{\text{avg}}^2 = \frac{1}{2} (\sigma_{v,x}^2 + \sigma_{v,y}^2). \quad (4.8)$$

By fitting the functional form of $\sigma_{\text{avg}}(t)$ to the experimental data, we can extract the temperature of the MOT. Such analysis is applied in Sections 4.1.1 and 4.1.2




Figure 4.2: TOF absorption images of the cloud released from the red MOT, taken at increasing TOF from 0 to 9 ms. The growth of the cloud size is fitted to extract the MOT temperature.

4.1.1 First Stage of Trapping: Blue MOT

The initial stage of atomic trapping is performed by the blue Magneto-Optical Trap (MOT), which utilizes the $^1S_0 \rightarrow ^1P_1$ transition at a wavelength of 461 nm. Before the atoms can be efficiently loaded into this trap, the atomic beam, which originates from the oven, must first be prepared.

Oven Operation and Source Lifetime

The atomic oven is maintained at a temperature of 420°C (Sec. 3.1.1). This temperature was selected as an optimal balance: it provides a sufficiently high atomic flux to ensure robust MOT loading while simultaneously guaranteeing a long operational lifetime for the source. Our estimations

suggest that at this temperature, the atomic reservoir can sustain daily use for approximately ten years before depletion. Although higher oven temperatures would increase the atomic flux (and thus the final MOT population), the trade-off would be a significant reduction in the source's lifetime.

Atomic Beam Preparation

To achieve efficient trapping, the atomic beam is decelerated using two primary components integrated into the atomic source (as detailed in Sec. 3.1.1): **Zeeman Slower**, which provides the necessary longitudinal slowing; **2D MOT**, which simultaneously performs transverse cooling and deflects the atomic beam toward the final trapping region.

Optimization of the Blue 3D MOT

The performance of the blue 3D MOT, situated within the glass cell, was meticulously optimized by scanning key experimental parameters. This optimization focused on maximizing both the total number of trapped atoms and the peak atomic density. The varied parameters included:

- the quadrupole magnetic field gradient,
- the power and detuning of the 3D MOT beams,
- the power and detuning of the 2D MOT beams,
- the power and detuning of the Zeeman slower,

The detailed results of these parameter scans are as follows.

Optimal Parameters

The following optimal parameters were determined based on system constraints and the goal of maximizing performance:

- **2D MOT Beams:** Powers were set to 9.4 mW ($s = 0.11$) and 4.7 mW ($s = 0.055$) for the two paths, with a common detuning of -22 MHz ($-0.7 \Gamma_{1P_1}/2\pi$).
- **Zeeman Slower:** Optimal power was 40 mW ($s = 0.93$) with a significant detuning of $-19 \Gamma_{1P_1}/2\pi = -580$ MHz.

- **3D MOT Beams:** Optimal power was 4.3 mW per beam ($s = 0.1$) with a detuning of $-2\Gamma_{1P_1}/2\pi = -60$ MHz relative to the atomic resonance.
- **Magnetic Field Gradient:** The best compromise between atom number and density was achieved with a gradient of 50 G cm^{-1} .

Enhanced Trapping via Shielding Effect

To further boost the number of trapped atoms, we implemented the **Shielding effect** [120]. This technique involves resonantly driving the red $^1S_0 \rightarrow ^3P_1$ intercombination transition (689 nm) concurrently with the blue MOT loading phase. In a highly saturated regime, the 689 nm light effectively shelves atoms in the $^3P_1$ state for approximately half of the time, thereby reducing the effective atomic density in the ground and excited $^1P_1$ states by up to a factor of two.




Figure 4.3: Relative enhancement of the atom number by the red-shielding beam as a function of detuning from resonance, with the off-resonant values taken as the baseline. Three distinct peaks are observed, corresponding to the magnetic sublevels of the $^3P_1$ state, which split into three energy levels in the presence of a magnetic field. The peaks are located at -11.0 MHz, -1.2 MHz, and +8.1 MHz, with an uncertainty of ± 0.5 MHz obtained from the gaussian fit and bootstrap analysis. The positions are approximately symmetric.

This density reduction suppresses light-assisted collisional losses during loading. When the red-shielding light is switched off, the atoms rapidly decay back to the $^1S_0$ ground state, restoring the full density and resulting in a final atom number that is approximately doubled compared to the unshielded case. Overall, this mechanism yielded a remarkable 80% improvement in the total atom number relative to the unshielded configuration, as illustrated in Fig. 4.3.

Addressing Loss Channels with Repumping Lasers

Despite the $^1S_0 \rightarrow ^1P_1$ transition being nearly closed, a critical loss pathway exists: atoms in the excited $^1P_1$ state can decay into the $^1D_2$ level, which subsequently populates the long-lived $^3P_2$ dark state. Atoms shelved in this state are removed from the main cooling cycle, severely degrading the MOT loading efficiency. As already discussed previously in Sec. 3.1.3, prevent the loss of these atoms, we employ two crucial repumping lasers:

- **707 nm Laser:** it couples the $^3P_2$ state to the $^3S_1$ state. From the $^3S_1$ state, atoms can decay into any of the three ^3P_J states;
- **679 nm Laser:** it repumps population specifically out of the $^3P_0$ state.

Together, these lasers ensure that the atoms are eventually returned to the ground state, allowing them to re-enter the cooling cycle.

As shown in Fig. 4.4, the impact of these repumpers is substantial:

- Without repumping, the MOT saturates quickly to a density $\sim 10^9$ atoms cm^{-3} , corresponding to less than 1×10^6 atoms.
- Using only the 707 nm repumper partially increases the atom density by a factor of 2, but introduces $^3P_0$ as a new decay channel from $^3S_1$.
- The simultaneous application of both the 707 nm and 679 nm repumpers eliminates the major loss pathways, boosting the density by a factor of 4 and the atom number by nearly an order of magnitude to 1.2×10^7 atoms.




Figure 4.4: Number of atoms in the blue MOT of ^{88}Sr as a function of loading time for different repumping configurations in absence of red-shielding light. The simultaneous use of 707 nm and 679 nm repumpers eliminates the major loss channels, leading to the highest observed atom number.

Final MOT Performance and Temperature

The temperature of the MOT was determined by monitoring the free expansion of the atomic cloud via the time-of-flight (TOF) method. The average cloud radius (σ_{avg}) was measured as a function of TOF (see Fig. 4.5), and fitting this data to the expansion model yielded a temperature of $T = 6.4 \text{ mK}$, which is consistent with published literature.

Summary of Optimized Blue MOT

Under the fully optimized conditions, incorporating both repumpers and the shielding technique, the first-stage blue MOT achieves the following performance metrics:

- **Trapped Atom Number:** 1.25×10^7 atoms
- **Peak Atomic Density:** $4 \times 10^9 \text{ cm}^{-3}$
- **Temperature:** 6.4 mK

These robust results successfully establish the foundation for the subsequent stages of cooling and trapping.




Figure 4.5: Average cloud radius σ_{avg} as a function of Time-of-Flight (TOF). The experimental data is fitted with the expansion model (solid line) to determine the MOT temperature as discussed in Eq. 4.7.

4.1.2 Bridging Stages: Broad-Band Red MOT

The final stage of cooling utilizes the narrow $^1S_0 \rightarrow ^3P_1$ intercombination transition (the "red" transition) to achieve microkelvin temperatures. Since this transition has an intrinsically low capture velocity, two techniques are implemented sequentially: the **Broad-Band MOT (BB MOT)** for efficient capture, followed by the **Single-Frequency MOT (SF MOT)** for ultimate cooling.

The BB MOT technique is commonly employed to increase the capture velocity, facilitating the transfer of atoms from the blue MOT. This method relies on creating a multi-frequency laser spectrum to broaden the effective cooling range. A recent alternative is the Sawtooth-Wave Adiabatic-Passage (SWAP) MOT.

Implementation and Frequency Comb

- **Waveform Generator:** The RF signal driving the AOM for the red 3D MOT beams is sinusoidally modulated by an Arbitrary Waveform Generator (Siglent SDG6022X AWG), which produces a carrier frequency sinusoidally modulated at 30 kHz with a modulation depth of ± 2.4 MHz.

- **Spectral Output:** The double-pass AOM configuration results in a frequency comb 9.6 MHz wide, with 30 kHz spacing between spectral components.
- **Detuning:** All frequency components are kept **red-detuned**. The upper edge is positioned just below the $^1S_0 \rightarrow ^3P_1$ resonance, ensuring that faster atoms interact with the more-detuned components, while atoms progressively cooled become insensitive to the least red-detuned frequencies.

Optical Parameters and Performance

- **Power Distribution:** Each 3D MOT beam delivers 4 mW of power, distributed over 320 frequency components.
- **Saturation Parameter:** Given a beam waist of 8 mm and $I_{\text{sat}} = 3 \mu\text{W cm}^{-2}$, the saturation parameter per tone is $s \approx 4.2$.
- **Power-Broadened Linewidth:** This corresponds to a power-broadened linewidth of $\Gamma_s/2\pi = 17 \text{ kHz}$.
- **Temperature:** The measured temperature of the red MOT during the broadband stage is $T = 21 \mu\text{K}$, ensuring a nearly uniform spectrum extending over all comb width.

Reducing the Magnetic Field Gradient

A rapid reduction of the quadrupole magnetic field gradient is crucial for trapping atoms using the narrow-line red transition.

- **Required Change:** The gradient must be reduced from the 50 G cm^{-1} required for the blue MOT to just a few G cm^{-1} .
- **Switching Protocol:** The MOT coils (operating at their peak of 100 A during blue MOT) are switched off completely, exploiting the fast switch-off circuit to decay the current in less than 300 μs . They are then re-enabled at a lower current of 4 A.
- **Transient Mitigation:** To suppress the transient overshoot and damped oscillations that occur upon coil reactivation, a voltage limit is implemented on the power supply. This helps reduce, but not completely eliminate, the overshoot caused by competition between the voltage and current control loops (see Fig. 4.6).

- **Optimal Gradient:** A final gradient of 2 G cm^{-1} was found to maximize the number of trapped atoms in the red MOT.




Figure 4.6: Measured magnetic field gradient during the transition from the blue MOT to the red MOT. The gradient is rapidly reduced from 50 G cm^{-1} to 2 G cm^{-1} by switching off and reactivating the coils at a lower current, simultaneously with turning off the blue MOT light and turning on the red MOT light.

4.1.3 Second Stage of Trapping: Single-Frequency Red MOT

The SF MOT is the final stage, designed to achieve the lowest temperatures and highest densities, following the BB MOT which maximized capture efficiency.

BB to SF Switching

A sudden switch is implemented, as opposed to a gradual spectral narrowing, for experimental simplicity:

- **Power Broadening:** A single-frequency component with a high saturation parameter of $s \approx 200$ is used. This power broadening effectively bridges the transition from the multi-frequency to the single-frequency regime.
- **SF Linewidth:** The resulting power-broadened linewidth is:

$$\Gamma_s/2\pi = (\Gamma/2\pi)\sqrt{1+s} \approx 105 \text{ kHz}$$

- **Switch Mechanism:** The RF source is switched from the Siglent AWG to a fixed-frequency RF channel of the Moglabs QRF driver via a TTL-controlled Minicircuits switch (ZX80-DR230-S+).
- **Efficiency:** This step achieves a high transfer efficiency, exceeding 80% from the BB MOT to the SF MOT.

Optimization and Final Performance

Measurements were taken at the end of a fixed 50 ms SF MOT stage, following a variable BB MOT loading time:

- **Optimal Loading Time:** Both atom number (N) and peak density (n_{peak}) saturate after approximately 250 ms of BB MOT loading (see Fig. 4.7), indicating that steady-state conditions are reached.
- **Peak Density:** The SF red MOT achieves the required high density of $n_{\text{peak}} \approx 2.5 \times 10^{10} \text{ cm}^{-3}$, essential for efficient loading into optical tweezers.
- **Final Atom Number:** $N \approx 5 \times 10^5$ atoms.




Figure 4.7: Atom number (left) and peak density (right) in the single-frequency red MOT as a function of the broadband MOT loading time. Saturation occurs around 250 ms.

Final Temperature

- **Comparison to Doppler Limit:** The theoretical Doppler cooling limit (T_D) for the power-broadened transition ($s \approx 200$) is $T_D = \hbar\Gamma_s/2k_B \approx 5 \mu\text{K}$. The achieved temperature is in good agreement with the theoretical minimum.

- **Measured Temperature:** TOF measurements yield a final temperature of $5.4(5) \mu\text{K}$ (see Fig. 4.8).

The entire sequence successfully transforms the atomic distribution from a large, dilute ensemble to a small, dense cloud, making it suitable for subsequent experiments.




Figure 4.8: Time-of-flight measurement for ^{88}Sr atoms in the single-frequency MOT. The average cloud waist is shown as a function of time of flight (dot with error bars), together with the fit (black line) used to extract the temperature. From the fit we obtain a temperature of $5.4(5) \mu\text{K}$.




Figure 4.9: Absorption images of the Strontium MOT at different operating stages. From left to right: broad-line shielded blue MOT, broadband red MOT stage, and narrow-line red MOT. The reduction of the cloud size and the increase in density are clearly visible as the cooling sequence progresses.

4.2 SINGLE ATOMS IN OPTICAL TWEEZERS

In the previous section, we introduced the principles and practical implementation of magneto-optical trapping as a means to prepare cold atomic ensembles. Although a MOT provides large numbers of ultracold atoms with relatively low experimental overhead, the confinement it offers is fundamentally collective: atoms remain weakly bound within an extended cloud and continue to interact with one another and with the cooling light.

The next step is to isolate and control individual atoms using optical tweezers. Transferring atoms from a MOT into an optical tweezer is far from straightforward. It requires careful engineering of the spatial and temporal overlap between MOT and tweezer, optimization of cooling and capture conditions, and control of heating processes and collisional losses. Once an atom is loaded into the tweezer, it must be further cooled and reliably detected while preserving both its survival probability and the fidelity of the measurement.

This section presents the full sequence used to achieve single-atom trapping in optical tweezers. We begin by examining the loading dynamics from the MOT, then describe the properties of the optical potential itself. We subsequently discuss the collisional processes and cooling mechanisms that enable single-atom occupancy, and we conclude with the imaging techniques employed to detect individual atoms and the methods used to characterize the system in terms of detection fidelity, survival probability, and atom temperature.

4.2.1 *Tweezer generation*

The optical setup used to generate the tweezer array is the one described in Sec. 3.2.4. In the current implementation, only the module based on the two crossed AODs has been realized and is actively used, while the SLM branch is still under testing. The crossed AODs allow us to generate static rectangular arrays. Although this configuration does not offer the same flexibility as an SLM, it nevertheless enables the creation of stable optical tweezers suitable for trapping single atoms.

To verify correct projection of the optical tweezers, a second “service” objective, identical to the main one, is mounted on the opposite side of the glass cell. This objective forms an imaging system coupled to a Basler Ace acA4024-29um camera, allowing direct imaging of the tweezer array (Fig. 4.10) and evaluation of optical aberrations.

Fine adjustments of the tweezer injection through the main objective are performed to minimize distortions. A precise parallel and coaxial alignment of the two objectives is essential, as any misalignment of the service imaging path could introduce additional aberrations that would distort the observed tweezer profiles. A simple and effective alignment procedure consists in first injecting a single static tweezer and imaging it through the service objective while iteratively adjusting tip-tilt and transverse position of the main objective. The alignment is optimized by minimizing asymmetries and ellipticity in the observed spot, as well as ensuring that the tweezer position remains invariant when slightly translating the focus along the optical axis.




Figure 4.10: Image of the optical tweezer array recorded using the service imaging system. The array consists of 3×3 traps generated with crossed AODs. Each tweezer has a waist of $\sim 1.5 \mu\text{m}$ and a center-to-center spacing of $10 \mu\text{m}$.

All results presented in this chapter were obtained using a 3×3 array, shown in Fig. 4.10. Each tweezer has a waist of approximately $1.5 \mu\text{m}$ and a center-to-center spacing of $10 \mu\text{m}$. The waist was experimentally determined by direct imaging of a single tweezer through the service objective, extracting its intensity profile at focus. The measured value exceeds the expected diffraction-limited waist ($0.9 \mu\text{m}$), most likely due to a slight underfilling of the objective's back focal plane and residual optical aberrations in the current setup. The extracted waist is consistent with the independently measured AC Stark shift inferred from the displacement

of the Sisyphus cooling resonance, providing spectroscopic validation of the optical calibration.

Although larger arrays with smaller spacings could in principle be produced, the present configuration is sufficient for the initial characterization of the system and for single-atom trapping. The choice of a small array is motivated not by technical limitations (e.g. FOV or available laser power), rather by the desire to simplify equalization of the trap depths across the tweezers. Future experiments will easily reach 10×10 arrays and beyond.

Loading atoms into optical tweezers requires precise spatial overlap between the tweezer array and the red MOT. This task is challenging because the field of view of the high-NA objective is only a few hundred micrometers, comparable to the size of the MOT itself ($\approx 400 \mu\text{m}$). To facilitate alignment, we use a 2 mm waist reference blue beam intersecting the red MOT orthogonally to the cell. This beam serves as a guide for positioning the microscope objective on its translation stage (manual x , y and motorized z). The axial position of the objective is then fine-tuned by recording fluorescence images of the red MOT collected through the objective. Thanks to the sub-micrometer repeatability (200 nm) of the axial translation stage, we can scan the MOT across different focal planes with high precision.

The overlap between the MOT and the tweezer array is optimized by displacing the MOT using bias magnetic fields along the three Cartesian axes. These fields are provided by the compensation coils described in Sec. 3.1.2, with a calibrated displacement sensitivity of $2 \mu\text{m}/(\text{mG}/\text{cm})$. This method is more robust than shifting the tweezer array itself, as the optical traps are extremely sensitive to alignment and aberrations. The overlap is continuously monitored using the service imaging system, which simultaneously records both the MOT fluorescence and the tweezer pattern.

Using this procedure, we achieve a good spatial overlap between the MOT and the tweezer array, enabling the first detection of fluorescence from atoms loaded into the optical traps. Subsequent fine adjustments of both alignment and experimental parameters are performed by directly maximizing the fluorescence signal from the trapped atoms. Precise, site-resolved imaging of single atoms is a key capability in modern quantum simulation platforms based on optical tweezer arrays. In our experiment, we employ fluorescence imaging to detect individual atoms, a technique that has become a standard toolbox element in neutral-atom experiments [121, 122, 123, 124].

The imaging sequence has two primary objectives:

- to determine whether an atom is present in a given tweezer with high reliability,
- to do so in a non-destructive manner, ensuring that the atom remains trapped for subsequent experimental steps.

The first requirement is quantified by the *detection fidelity*, which measures the ability to discriminate atomic fluorescence from background noise. The second is characterized by the *survival probability*, which quantifies the likelihood that an atom remains in the trap following the imaging exposure.

Both quantities are strongly linked to the photon-scattering process. Sufficient fluorescence must be collected to distinguish signal from background, while at the same time mechanisms that degrade detection must be limited, namely: heating induced by photon recoil, which may eject the atom from the trap, and optical pumping into long-lived dark states that interrupt the imaging cycle. Thus, imaging requires not only illumination with a resonant beam but also continuous cooling and repumping to maintain a closed scattering cycle.

High-fidelity detection and high survival probability are essential for the realization of large-scale, defect-free arrays and for enabling controlled interactions for many-body quantum simulation.

4.2.2 Imaging scheme

Two imaging approaches have been demonstrated for strontium atoms. The first uses the strong dipole-allowed $^1S_0 \rightarrow ^1P_1$ transition at 461 nm [68, 69, 125], while the second exploits the narrow $^1S_0 \rightarrow ^3P_1$ intercombination line at 689 nm [55]. Both enable single-atom detection, but with different trade-offs in scattering rate, heating, and technical complexity. In the present setup, we employ the 461 nm scheme, while the 689 nm scheme remains an attractive path for future non-destructive imaging.

For better understanding we invite the reader to take a look at Fig. 2.6, where the electronic structure of ^{88}Sr is shown. The 461 nm transition has a large natural linewidth ($\Gamma/2\pi \approx 32$ MHz), providing bright and fast imaging. However, it is not perfectly closed: a small branching ratio from $^1P_1$ to $^1D_2$ leads to population in the metastable ^3P_J states. Without repumping, atoms would leave the imaging cycle.

We therefore employ the same repumpers used in the blue MOT to return population from $^3P_2$ and $^3P_0$ back to the ground state.

To counteract heating during imaging in the optical tweezer traps, we apply simultaneous cooling on the 689 nm transition using the Sisyphus cooling scheme described in Sec. 2.1.7. Because atoms populate the $^3P_1$ state during cooling, blue scattering is intermittent. To avoid excessive heating and maximize survival, imaging is performed using short pulses of blue light interleaved with continuous red cooling. This approach allows efficient fluorescence collection while minimizing heating, as the red light AC Stark shift remains negligible.

The trapping light at 813.4 nm induces a differential light shift between the $^1S_0$ and $^1P_1$ states. In a typical tweezer (trap depth ~ 360 μ K), the excited state is more strongly confined, leading to a red shift of about -12 MHz. The imaging beam detuning must therefore be adjusted. Indeed, one can even detune the imaging beam such that resonance occurs near the trap center, naturally reducing the scattering rate toward the edges and improving survival.

The alternative red imaging scheme, based on scattering from the 7.6 kHz-wide intercombination line, enables nearly non-destructive imaging [55]. Because heating is negligible, atoms can remain near the motional ground state throughout imaging without the need for repumpers. However, this scheme demands precise control over detuning and intensity, and imaging must often be interleaved with cooling at millisecond duty cycles. The longer wavelength also limits spatial resolution and increases imaging time.

4.2.3 Photon collection

To detect single atoms, we collect fluorescence emitted from atoms illuminated by a resonant imaging beam incident orthogonally to the objective. Because each photon contributes to distinguishing atomic signal from background, maximizing collection efficiency is crucial.

The fluorescence imaging system is shown schematically in Fig. 4.11. A resonant 461 nm imaging beam along $\hat{z}$ excites the atoms, while the 689 nm cooling beam is overlapped along the same axis. Fluorescence is collected by the same high-numerical-aperture objective used for tweezer generation (Sec. 3.2.1). A dichroic mirror separates the tweezer light from the atomic fluorescence, which is relayed to the detection camera.

The polarization of the imaging beam determines the atomic dipole radiation pattern and strongly affects collection efficiency.




Figure 4.11: Fluorescence imaging setup. The tweezer beam (red) is focused into the cell by a high-NA objective. A resonant imaging beam along $\hat{z}$ excites the atoms, and Sisyphus cooling is provided by the 689 nm beam. Fluorescence collected by the same objective is separated from the tweezer light by a dichroic mirror and imaged onto a low-noise camera.

Following the formalism of Ref. [52], the dipole operator is decomposed into spherical components d_q ($q = 0, \pm 1$), corresponding to π and $\sigma^\pm$ transitions. For our $J_g = 0 \rightarrow J_e = 1$ transition, only $q = -m_e$ contributes. The radiation patterns simplify to [60]:

$$D(\theta, \phi) = \frac{3}{8\pi} \sin^2 \theta \quad (\text{for } \pi \text{ polarization}), \quad (4.9)$$

$$D(\theta, \phi) = \frac{3}{16\pi} (1 + \cos^2 \theta) \quad (\text{for } \sigma^\pm \text{ polarization}), \quad (4.10)$$

where θ is measured with respect to the quantization axis $\hat{z}$.

In our configuration, the imaging beam propagates along $\hat{z}$ with a σ^+ polarization along $\hat{y}$, and the objective collects along $\hat{x}$, maximizing collected fluorescence. By integrating the dipole radiation pattern over a circular aperture, we extract the geometric collection efficiency for a given numerical aperture (NA) [126]:

$$C = \frac{1}{8} \left[4 + (\text{NA}^2 - 4) \sqrt{1 - \text{NA}^2} \right]. \quad (4.11)$$

For our objective ($\text{NA} = 0.56$), we obtain $C = 0.118$. Future upgrades will use the second objective to collect fluorescence from the opposite side and retroreflect it, effectively doubling the solid angle.

A crucial component of the imaging system is the detection camera. We use a scientific qCMOS camera (Hamamatsu ORCA-Quest C15550-20UP), which combines high quantum efficiency ($\sim 85\%$ at 461 nm, $> 55\%$ at 689 nm) with ultra-low readout noise ($< 0.9 e^-$ RMS). Unlike CCDs, the CMOS architecture enables fast parallel readout, large sensor formats (4096×2304 pixels of $4.6 \mu\text{m}$), and flexible trade-offs between speed and sensitivity. The low readout noise enables resolving photon-number peaks with $< 1\%$ misclassification error, essential for quantitative atom detection.

The fluorescence is reimaged onto the sensor via a tube lens of focal length $f_{\text{tl}} = 100 \text{ mm}$, yielding a magnification $M = f_{\text{tl}}/f_{\text{EFL}} \approx 4.2$ with $f_{\text{EFL}} = 24 \text{ mm}$ the effective focal length of the objective. With this magnification, the image of a single atom spans roughly one camera pixel, maximizing signal contrast.

Residual aberrations enlarge the point spread function (PSF) beyond the diffraction-limited size. These are minimized by careful alignment and the use of aberration-corrected optics. Background light from the tweezer path is suppressed using a dichroic mirror and spurious reflections from the tweezer path are suppressed by three bandpass filters.

Overall, the combination of a high-performance qCMOS detector, high-NA optics, and aggressive filtering yields an imaging system capable of resolving single atoms with high fidelity, while remaining compatible with future extensions to fully non-destructive red imaging.

4.2.4 *From multiple to single atoms*

Once a first fluorescence signal from atoms confined in the tweezers was obtained, the experimental parameters were further optimized to enhance the detected fluorescence. An example is shown in Fig. 4.12, where multiple tweezers are simultaneously loaded with many atoms. In this particular realization, a 3×3 array of tweezers is employed, each with a trap depth of approximately $360 \mu\text{K}$.

In order to maximize the fluorescence signal, it is essential to reduce atom loss during the imaging stage. To this end, cooling is applied to the trapped ensemble, ensuring that the atoms remain confined while scattering photons. The fluorescence yield thus serves not only as an




Figure 4.12: Fluorescence image demonstrating the simultaneous loading and detection of multiple atoms in a 3×3 array of optical tweezers, with a trap depth of approximately $360 \mu\text{K}$ per tweezer. The signal shows distinct trapping sites, with the color map indicating the detected photon counts (intensity). The axes represent the spatial extent of the imaging region in micrometers.

indicator of the atom number, but also as a sensitive probe of the efficiency of the cooling stage.

In addition to the measurements averaged over the 9 tweezers, individual tweezer signals were analyzed. This allowed for fine-tuning of the trap depths across the array. Specifically, by adjusting the amplitudes of the RF tones that generate the individual traps via AOD diffraction, the depths could be equalized to minimize inhomogeneities among the tweezers. Such balancing is crucial: small differences in trap depth can lead to different Stark shifts for the 461 nm and 689 nm transitions, resulting in variations in loading probability, cooling efficiency, and fluorescence across the array. By carefully optimizing these parameters, a more uniform array of traps is achieved.

For fluorescence imaging, we employed a sequence of pulsed blue light with a total duration of 250 ms, consisting of 1 ms pulses repeated at a frequency of 1 kHz (1 ms pulses at 333 Hz repetition rate, $s = 0.03$). During the imaging sequence, a single beam implementing Sisyphus cooling was kept on, ensuring that the atoms remained confined in the tweezers while scattering photons.




Figure 4.13: Two-dimensional map of fluorescence yield during Sisyphus cooling optimization. The photon count (color scale) is plotted as a function of both the cooling beam detuning (x-axis) and the saturation parameter S (y-axis). The distinct bright region highlights the optimal operational window where maximum fluorescence is achieved, providing a clear map for simultaneously setting the ideal cooling frequency and intensity.

To identify the optimal cooling parameters, systematic scans of the Sisyphus cooling beam frequency and power were performed. The polarization of the Sisyphus beam was also empirically optimized to maximize the resulting signal. The results of these scans, shown in Fig. 4.13, represent averages over the 9 tweezers in the array and provide a clear picture of the parameter regions where the fluorescence is maximized. Based on these optimization scans, we operated the Sisyphus cooling beam at a detuning of -1.4 MHz relative to resonance and at a saturation parameter of approximately 5×10^3 .

The intensity of the imaging beam was also optimized in order to maximize the detected fluorescence signal, as shown in Fig. 4.14. The fluorescence signal is observed to increase sharply at low imaging powers and then to saturate, with an optimal working point at $s = 0.1$. Operating in this regime provides a good compromise between achieving high photon counts and avoiding excessive perturbation of the trapped atoms.




Figure 4.14: Fluorescence counts from multiple atoms in tweezers as a function of the imaging beam saturation parameter s . The signal rises steeply at low power and then reaches a plateau, indicating the onset of saturation. The optimal working point was chosen around $s = 0.1$, balancing photon collection efficiency with minimal heating and atom loss.

It should be emphasized that fluorescence images obtained under these conditions do not allow us to determine the number of atoms trapped in each tweezer. The strong fluorescence signal indicates multiple-atom occupancy, but the detection does not provide the single-atom sensitivity required to distinguish between different atom numbers. These measurements therefore serve primarily as a diagnostic for optimizing trap loading and imaging conditions.

Single atoms

In order to reach the single-atom regime, a light-assisted collisions (LAC) stage is introduced in the sequence. This procedure, described in Sec. 2.1.6, removes multiple occupancies and ensures that each tweezer is either empty or contains a single atom [59]. This additional step is implemented immediately after the initial loading, using parameters similar to those employed for fluorescence imaging [127]. We scanned the Sisyphus beam frequency and power, the imaging beam power, and the LAC duration to identify the operating conditions that produce a clear bimodal separation in the fluorescence count distribution (associated with 0 or 1 atom)

while simultaneously maximizing the atom survival probability. Under optimized conditions, we typically achieve a filling fraction of about 50%, consistent with the probabilistic nature of single-atom loading.

After this stage, the system enters the single-atom regime. Due to the stochastic nature of the collisions that leave at most single trapped atoms, not all sites are occupied in a given realization, as illustrated in Fig. 4.15. To visualize the array more clearly, we average over 1000 realizations, shown in Fig. 4.16(a). A complementary view is provided by the fluorescence count histogram in Fig. 4.16(b), constructed from the same dataset and analyzed as described in Sec. 4.2.5. The histogram exhibits a well-defined bimodal distribution with two clearly separated peaks corresponding to empty (0 atoms) and occupied (1 atom) tweezers. The absence of intermediate or higher-count peaks demonstrates that multiple occupancies are effectively suppressed by the LAC process.




Figure 4.15: Fluorescence images of the 3×3 tweezer array in the single-atom regime, taken in consecutive experimental realizations. Due to the stochastic nature of the LAC process, not all sites are occupied in each shot.

With these optimized conditions, the LAC procedure reliably prepares arrays of single atoms, enabling subsequent experiments that require deterministic single-particle control. The optimal LAC parameters are:

$$s_{461} = 0.15, \quad s_{689} = 2 \times 10^4, \quad \Delta_{689} = -2 \text{ MHz},$$

where Δ_{689} denotes the detuning relative to the free-space resonance of the 689 nm transition.

It is worth noting that the optimal Sisyphus cooling parameters and imaging powers for detecting single atoms differ from those used in the many-atom regime. After adjusting the parameters to balance photon collection efficiency with minimal heating and to ensure high-fidelity single-atom detection, we find:

$$s_{461} = 0.03, \quad s_{689} = 1.3 \times 10^4, \quad \Delta_{689} = -2.1 \text{ MHz}.$$




Figure 4.16: (a) Fluorescence image of a 3×3 tweezer array in the single-atom regime, obtained by averaging over 1000 experimental realizations after the LAC process. Here, photo counts are pixel-relative. (b) Histogram of fluorescence counts from the same dataset, showing a bimodal distribution with two clearly separated peaks corresponding to empty and singly occupied tweezers. The absence of additional peaks confirms suppression of multiple-atom occupancies. Here, photo counts are ROI-relative (see Sec. 4.2.5).

The full experimental sequence used to prepare single atoms is summarized in Fig. 4.17. The protocol begins with a blue shielded MOT stage, providing initial cooling and compression. This is followed by a multifrequency red MOT stage that further reduces the temperature while maintaining a high atomic number. A subsequent single-frequency red MOT enables efficient transfer of atoms into the optical tweezer array, operated at a trap depth of $360 \mu\text{K}$.

We note that the durations of the blue MOT and red BB MOT stages were reduced from their optimal standalone values to keep the sequence as short as possible for the purpose of accumulating statistics, while still maintaining sufficiently efficient loading into the tweezers.

Immediately after loading, a LAC sequence for 30 ms is performed by shining a $\Delta_{\text{LAC},689}/2\pi = -2 \text{ MHz}$ red-detuned 689 nm beam ($s = 2 \times 10^4$) and we speed up the loss process by adding pulsed resonant 461 nm light with a 1ms duration and 333 Hz repetition rate ($s = 0.15$), as in [43]. Although the LAC and imaging stages share a similar structure, their parameters (beam powers, detunings, durations) are independently optimized: the LAC step maximizes single-atom preparation fidelity, while the imaging step ensures efficient photon collection without compromising atom survival.




Figure 4.17: Experimental sequence for atom loading, preparation, and imaging. The timeline shows the successive stages: blue MOT, multifrequency red MOT, single-frequency red MOT, LAC, and imaging. Plotted as functions of time are the intensities of the MOT beams, magnetic field gradient, tweezer depth, Sisyphus beam power, and imaging beam power. During the final imaging stage, 1 ms blue-light pulses at 333 Hz are applied for 250 ms, together with continuous Sisyphus cooling. This sequence ensures efficient loading, elimination of multiple occupancies, and high-fidelity single-atom fluorescence detection.

4.2.5 Signal analysis

The fluorescence signal collected by the camera, as described in Sec. 4.2.3, is analyzed to determine whether an atom is present in a given optical tweezer. In the regime discussed in Sec. 4.2.4, the outcome of this analysis is inherently binary: an atom is either detected or not. To reliably distinguish atomic fluorescence from background counts, a sufficient number of photons must be collected during the exposure.

The first step in the analysis is the accurate identification of the atomic positions within the acquired fluorescence images. For this purpose, a large dataset of single-shot images is recorded and averaged to obtain a clean reference image of the tweezer array, as shown in Fig. 4.16. The resulting averaged image displays well-defined fluorescence maxima corresponding to the positions of the trapped atoms.

For each fluorescence peak, a region of interest (ROI) is defined around its center. Given the magnification of our imaging system, the PSF width is expected to span approximately one pixel. We therefore select ROIs of 3×3 pixels centered on each atom, as illustrated in Fig. 4.18.




Figure 4.18: Example fluorescence image showing the selected 3×3 pixel ROIs centered on the atomic array sites. These ROIs are used to extract the photon counts for each individual atom from a single acquisition.

The choice of ROI size is crucial: a too small region underestimates the signal, while a too large region introduces unnecessary background noise. Moreover, for closely spaced atoms, large ROIs may include fluorescence from neighboring tweezers. In the present work, the interatomic spacing is sufficiently large to avoid PSF overlap. In situations where adjacent PSFs are not fully resolved, principal component analysis and several deconvolution techniques have been demonstrated to provide effective disambiguation [128].

The background level is estimated by sampling image regions that contain no tweezers. The mean background value is subtracted from the entire image to isolate the atomic fluorescence signal. The camera output, expressed in Analog-to-Digital Units (ADU), is then converted into electrons using the system gain factor (0.11 electrons per ADU for our camera). Because not all incident photons generate electrons, this signal is further corrected using the quantum efficiency (QE) of the sensor. In our case, the QE is 0.85, meaning that 85% of the incoming photons

produce a detectable electron. Summing the calibrated signal within each ROI yields the photon count associated with each tweezer.

Finally, the presence or absence of an atom is determined by comparing the photon count to a predefined discrimination threshold. Counts below this threshold correspond to empty tweezers, whereas higher counts indicate the presence of an atom. Although the threshold can in principle be chosen arbitrarily, in practice it is optimized to maximize the detection fidelity, as described in the following section.

Fidelity

After each fluorescence image is acquired, the presence or absence of an atom in each tweezer is classified according to the method described above. We define the *imaging fidelity* F as the probability that this classification is correct—that is, the probability that an atom is identified as present when it was indeed present at the start of the exposure, or identified as absent when no atom was initially present. This definition explicitly refers to the beginning of the exposure time, since atoms may be lost during the imaging process.

The number of detected photons follows a stochastic process governed by the statistics of spontaneous emission and photon detection. Ideally, the photon counts follow a Poisson distribution [129]. In practice, however, camera readout noise and background light-scattering broaden and distort this distribution. As a result, the experimentally measured histogram exhibits two peaks: one corresponding to empty tweezers (background) and another corresponding to occupied tweezers (atomic fluorescence), as shown in Fig. 4.19.

To model these two populations, we fit the histogram using a bimodal skewed Gaussian distribution. Each peak is described by the skewed Gaussian function [96]

$$S(x, A, \mu, \sigma, \alpha) = \frac{A}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(x - \mu)^2}{2\sigma^2}\right] \left[1 + \operatorname{erf}\left(\frac{\alpha(x - \mu)}{\sqrt{2}\sigma}\right)\right], \quad (4.12)$$

where A is the amplitude, μ and σ are the mean and standard deviation, and α controls the skewness. The overall model for the photon-count histogram is the sum of two such functions:

$$P(x) = S_0(x, A_0, \mu_0, \sigma_0, \alpha_0) + S_1(x, A_1, \mu_1, \sigma_1, \alpha_1), \quad (4.13)$$

representing the empty and occupied tweezers, respectively.




Figure 4.19: Photon-count histogram for a single ROI, acquired by accumulating for 100 different imaging sequences. The data are fitted with a bimodal skewed Gaussian model (red line), corresponding to the empty and occupied tweezers. The vertical dashed line indicates the optimal threshold (23 counts) used for atom classification, yielding a fidelity of 99.986(6)%.

From this fit, we determine the optimal detection threshold x_t that minimizes the total classification error,

$$\varepsilon(x_t) = \varepsilon_1(x_t) + \varepsilon_2(x_t), \quad (4.14)$$

where

$$\varepsilon_1(x_t) = \int_{x_t}^{\infty} S_0(x) dx \quad (\text{false positives}), \quad (4.15)$$

$$\varepsilon_2(x_t) = \int_{-\infty}^{x_t} S_1(x) dx \quad (\text{false negatives}). \quad (4.16)$$

The imaging fidelity is then

$$F = 1 - \varepsilon(x_t), \quad (4.17)$$

and the optimal threshold is obtained numerically by maximizing F .

Applying this procedure to our experimental data, we extract a threshold value of $x_t = 23$ counts and obtain a fitting cut-off fidelity of

$$F = 99.986(6)\%$$

for a trap depth of 360 μK and an imaging duration of 250 ms.




Figure 4.20: Measured imaging fidelity as a function of imaging duration for a fixed trap depth of $360 \mu\text{K}$. Longer exposure times improve photon statistics and increase the fidelity, which saturates near unity for exposure times above 250 ms.

The imaging duration may be reduced at the expense of fidelity, enabling faster duty cycles when required. Figure 4.20 shows the measured dependence of the fidelity on the exposure time, illustrating the trade-off between speed and accuracy. Uncertainty estimates for Fig. 4.20 and for subsequent results were obtained using the bootstrap method [130]. In this non-parametric approach, multiple synthetic datasets are generated by resampling the original data with replacement.

For each resampled dataset, the relevant statistical quantity is recalculated, and the resulting distribution provides an empirical estimate of its uncertainty. This method yields robust error bars without requiring assumptions about the underlying probability distributions.

Survival probability

The ability to maintain atomic confinement during fluorescence imaging is a crucial requirement in optical tweezer array experiments. Post-imaging rearrangement protocols and engineered interactions rely not only on identifying the atomic positions with high fidelity, but also on ensuring that atoms remain trapped throughout the imaging sequence. For this reason, in addition to optimizing the imaging fidelity F , it is essential to maximize the *survival probability*, defined as the likelihood that an atom present at the beginning of the imaging process remains in the trap at its end. We denote this quantity by S .

In the ideal limit of perfect imaging fidelity ($F = 1$), the survival probability could be directly obtained by acquiring two consecutive fluorescence images under identical conditions and measuring the fraction of atoms detected in both [131]. This fraction would correspond to S . Under realistic conditions, where $F < 1$, this method yields only an approximate estimate, since false positives (erroneous identification of empty traps as occupied) and false negatives (missed detection of atoms) introduce systematic biases. A more accurate determination of S can be obtained using models that explicitly account for these detection imperfections [60].

As a first-order approximation, and given the high fidelity achieved in our system, we employ the two-image method to estimate the survival probability and to investigate its dependence on the photon scattering rate during fluorescence imaging. Experimentally, we begin by preparing an array of single atoms in the tweezers, and we acquire a first fluorescence image to determine which traps are initially occupied. To prevent atom loss between successive exposures, both the trapping potential and the cooling light remain continuously active throughout the measurement, ensuring stable trapping conditions. A second image is acquired following a delay of 300 ms, corresponding to the camera readout time and the re-arming period required for the next trigger. This ensures that both images are collected under identical experimental conditions. By comparing the two images, we extract the survival probability as the fraction of atoms that remain trapped over the duration of the imaging sequence.

The resulting measurements are shown in Fig. 4.21. Each data point represents the mean survival probability obtained from repeated experimental runs, with error bars indicating the corresponding statistical uncertainty. As illustrated in the figure, the survival probability decreases sharply with increasing scattering rate during fluorescence imaging. At low scattering rates, atoms are almost always retained in the trap, whereas at higher rates photon-recoil heating becomes significant, leading to enhanced loss. A more detailed view of the data in the optimal scattering regime is presented in the insert of Fig. 4.21, where the survival probability is plotted together with the imaging fidelity. At sufficiently high scattering rates, atom loss during imaging prevents the formation of a clear bimodal fluorescence histogram, making the fidelity estimation less reliable.




Figure 4.21: Trapped atom survival probability and detection fidelity $\mathcal{F}$ (in the insert that zooms on the first points) as a function of the probe beam scattering rate. All data are taken with a fixed imaging duration of 150 ms, chosen to ensure that for every scattering rate the detected signal was sufficiently strong to allow accurate evaluation of both quantities. The shaded region indicates the optimal scattering-rate value. Error bars represent standard deviations calculated with the bootstrapping method. [130]

From this characterization, we identify an optimal operating point corresponding to a scattering rate of 3.12 MHz, which maximizes both survival probability and imaging fidelity. Under these conditions, we obtain a survival probability of $S = 97(2)\%$ and a fidelity of $F = 99.986(6)\%$.

Vacuum lifetime

The vacuum lifetime provides a direct and quantitative measure of the quality of the vacuum environment in which the atoms are confined. It is defined as the characteristic timescale over which a trapped atom is lost due to collisions with residual background gas atoms or molecules. Even under ultra-high-vacuum (UHV) conditions, a small but finite background gas density remains. Collisions between trapped atoms and these residual particles can transfer sufficient kinetic energy to expel the atoms from the optical potential, ultimately limiting the achievable trap lifetime.

A long vacuum lifetime is therefore a key figure of merit for experiments based on single trapped atoms. The loss of even a single atom may require reloading the tweezer, introducing dead time and reducing the overall experimental duty cycle. In quantum simulation experiments, the loss of one atom can modify the connectivity of the atomic array, resulting in systematic errors in the simulated dynamics. Ensuring long trap lifetimes is thus essential for carrying out complex experimental sequences without interruption and at high repetition rates.

To characterize the vacuum lifetime, single atoms are loaded into optical tweezers and an initial fluorescence image is acquired to identify the occupied traps. The trapping and cooling light are then left on for a variable hold time, t_{hold} , during which collisions with background gas particles may lead to atom loss. After this interval, a second fluorescence image is taken, and the fraction of atoms that remain trapped is extracted. By repeating this procedure for different values of t_{hold} , we obtain the survival probability $P(t)$ as a function of time.

The resulting measurements, shown in Figure 4.22, are fitted with a single-exponential decay model:

$$P(t) = e^{-t/\tau_{\text{vac}}}, \quad (4.18)$$

where τ_{vac} denotes the vacuum lifetime. From the fit, we obtain a vacuum lifetime of

$$\tau_{\text{vac}} = 425(25) \text{ s.}$$

This long lifetime confirms the high quality of the vacuum environment and demonstrates the suitability of the apparatus for performing extended and demanding experimental sequences with minimal atom loss.




Figure 4.22: Measurement of the vacuum lifetime. The survival probability of a single atom in the tweezer is plotted as a function of the hold time. The experimental data (black points with error bars) are fitted with an exponential decay function (gray line), yielding a lifetime of $\tau = 425(25)$ s.

Atom temperature

Another key parameter in experiments with individually trapped atoms is their temperature, which determines both the spatial and motional extent of the atom within the tweezer potential and is directly linked to the mean motional quantum number. It therefore serves as an essential diagnostic for evaluating trap quality and identifying possible sources of heating. Accurate knowledge of the temperature is crucial for optimizing subsequent steps such as the preparation of coherent quantum states and the implementation of high-fidelity operations in neutral-atom quantum-computing architectures. Furthermore, temperature measurements offer valuable insight into heating mechanisms, including trapping-laser intensity noise, mechanical vibrations, or residual photon scattering.

In our experiment, the atomic temperature was measured using the standard release-and-recapture method, a robust and widely adopted technique in atomic physics for characterizing the temperature of atoms confined in optical dipole traps. The principle is straightforward: a single atom initially held in the optical tweezer is released by switching off the trapping laser for a controlled duration t_{off} . The trap is then turned back on, and we record whether the atom is recaptured or lost.

During the free-expansion interval, the atom moves according to its initial thermal velocity. If its total energy remains below the trap depth, it stays within the capture region and is recaptured once the trap is restored; otherwise, it escapes.

Because the outcome of each realization is binary (recaptured or lost), the sequence is repeated many times for each value of t_{off} , yielding a statistical estimate of the recapture probability $P_{\text{rec}}(t_{\text{off}})$. The shape of this curve encodes the atomic velocity distribution and thus the temperature. However, an analytical relationship between temperature and recapture probability is not available, since the atomic motion in the Gaussian optical potential is nonlinear and influenced by gravity as well as by potential anisotropies. For this reason, the analysis relies on Monte-Carlo simulations that faithfully reproduce the full release-and-recapture process.

For the present work, we adapted the simulation framework developed in Ref. [132], which implements a classical thermodynamic model of the atomic trajectories. Atoms are initialized with positions and velocities sampled from a Boltzmann distribution at temperature T . Each trajectory is propagated during the release interval t_{off} in free space, including gravity, and the trap is then reactivated. An atom is deemed recaptured if its total energy remains below the trap depth. The model incorporates a realistic Gaussian potential, taking into account the dependence of the radial and axial trapping frequencies and trap depth on laser parameters such as wavelength, power, and waist. For each release time, the simulation performs a large number of independent realizations, and the resulting recapture fraction

$$P_{\text{sim}}(t_{\text{off}}, T) = \frac{N_{\text{recaptured}}}{N_{\text{total}}}$$

is compared with the experimental data. The temperature is extracted by minimizing the normalized deviation between simulation and experiment. Figure 4.23 shows the measured recapture probabilities together with the best-fit Monte Carlo simulation, which reproduces the data with excellent agreement. From this analysis, we extract an atomic temperature of $12.95(5)$ μK for atoms trapped in optical tweezers at 813 nm with a trap depth of 360 μK . No trap-depth ramping is applied, and Sisyphus cooling is implemented under optimal conditions, resulting in atoms prepared close to their motional ground state. This result confirms the effectiveness of the implemented cooling scheme and the high quality of the optical confinement.




Figure 4.23: Recapture probability as a function of the release time for release-and-recapture measurements. The experimental data (points with error bars) are compared with a Monte Carlo simulation (solid line). The extracted temperature is $12.95(5) \mu\text{K}$.

CASE STUDY: KINETICALLY-INDUCED FRUSTRATION DYNAMICS

This chapter is based on the work I carried out in my 4-months internship at Institut d'Optique in Prof Browaeys's group, and therefore related to the publication:

M Qiao, R Martin, L Homeier, I Morera, B Gély, L Klein, Y Chew, D Barredo, T Lahaye, E Demler, and A Browaeys:

"Kinetically-induced bound states in a frustrated Rydberg tweezer array", [arXiv:2510.17183](https://arxiv.org/abs/2510.17183) [physics.atom-ph] (2025).

During my third year, I joined the group of Prof. Antoine Browaeys at the Institut d'Optique in Palaiseau (Paris), for a compulsory internship, which took place from January 1st to April 30th, 2025. Over the course of this four-month collaboration, I worked within the CHADOQ laboratory, where a new experimental platform was being developed to study the "t-J-V Hubbard Model" in the context of *Magnon-Hole dynamics*. The goal of this project was to achieve first direct experimental observation of kinetically-induced bound states, where particles bind, not due to attractive interactions, but because joint motion reduces kinetic frustration.

My initial contribution focused on the finalization of the experimental platform, in particular through the implementation and calibration of spatial light modulator (SLM)-based holographic generation of optical tweezer arrays. The CHADOQ group adopted a new feedback algorithm for SLM mask creation, known as *Neighbor Weighted Gerchberg-Saxton* (NeWGS) algorithm (see Sec. 3.2.3). I carried out a systematic calibration of the optical response of the system and benchmarked the performance of this method, demonstrating an improvement in convergence speed of 66% (i.e. reducing iterations to 1/3). This resulted in a significant acceleration of the experimental preparation stage while maintaining high uniformity and reliability of the generated arrays.

I integrated this algorithm into the CHADOQ experimental workflow by implementing it within their custom graphical user interface for SLM control. Moreover, I improved the control software by leveraging prior experience gained within the Florence group, designing modular Python-based tools and Qt graphical interfaces to streamline experimental sequences, parameter management, and real-time diagnostics. All together these developments enhanced the overall stability, flexibility, and reproducibility of the experiment, enabling efficient exploration of the relevant parameter space of the "t-J-V Hubbard Model" Hamiltonian.

During the final month of my internship, I participated in the acquisition and analysis of the first experimental datasets. In particular, I focused on the precise calibration of the ladder geometry, where the inter-site distance plays a crucial role in determining tunneling amplitudes and interaction strengths, and thus the dynamical regime of the system. Establishing this calibration was important for reliably accessing the magnon-hole dynamics explored in the subsequent experimental campaign, conducted after my departure.

5.1 MAGNON-HOLE DYNAMICS

Quantum frustration can give rise to a rich variety of exotic many-body states, such as spin liquids and unconventional superconductors. For such a dynamics we can consider two component mixtures of fermionic or bosonic particles on lattices with frustrated kinetic energy, in the regime of light doping relative to half-filled insulating states, near-complete spin polarization, and repulsive SU(2) symmetric interactions. We can observe that these systems tend to form multi-particle bound states, where effective attraction emerges from kinetic frustration. In many cases we can find that the ground state of such systems is best understood as resulting from self-organization of multi-particle clusters.

The systems under consideration can be relevant for understanding twisted bilayer graphene close to ferromagnetic insulating states. We consider here one specific lattice geometry: **zigzag ladders**, but we expect that the conclusions may apply to broader classes of systems. The initial state features one particle per site on a frustrated ladder (either zigzag or flux-threaded) with full spin polarization. When one of the spins is flipped, we call it a *Magnon*. When a single particle is removed altogether, we call it a *Hole*. The starting point of our analysis is the observation of attraction between *holes* and *magnons* arising from frustration of the kinetic energy of holes.

Kinetic Frustration

While the most familiar mechanism for bound state formation involves attractive interactions between particles, kinetic frustration offers an alternative route. **Kinetic frustration** occurs in lattice systems when multiple quantum paths of a propagating particle interfere destructively. In such cases, the particle cannot fully exploit the kinetic energy available in a non-frustrated (e.g., unfrustrated lattice) system, due to cancellation between amplitudes along different paths.




Figure 5.1: Kinetic Frustration in a triangular plaquette. (a) Constructive interference of tunneling paths in a triangular geometry when all sites host identical atoms, leading to coherent amplitude enhancement. (b) The presence of a distinguishable atom (red) breaks path indistinguishability and suppresses interference, resulting in the absence of coherent enhancement. Arrows indicate tunneling paths and their relative phases.

This frustration can be alleviated by the presence of another distinguishable particle. Suppose the frustrated particle can move between two sites via two different paths. In one path, the second particle remains fixed; in the other, it is displaced to a different site. Because the final configurations of the two trajectories differ (due to the second particle's movement), the paths no longer interfere, and the frustrated particle can gain kinetic energy.

A single *Hole* injected into a fully polarized insulating state exhibits kinetic frustration arising from destructive interference of different paths. This occurs for non-bipartite geometries such as the zigzag ladder or for bipartite ones in the presence of magnetic flux.

When a *Magnon* is added into the system, the two different paths cannot interfere since they lead to different final states as shown in Figure 5.1. These two final states correspond to the permutation of the two spins. The primary goal of CHADOQ's work was to study different cases of kinetic frustration (n Holes, m Magnons) giving rise to formation of multi-particle bound states in a zig-zag ladder. In the experiment, bound states are identified not only by energetic separation in Energy-Dispersion (ED) spectra, but by persistent real-space hole-magnon correlations combined with delocalized center-of-mass motion.

The t -J Hubbard Model

The *Hubbard Model* plays a central role in the study of strongly correlated quantum systems. It describes particles, typically fermions or bosons, trapped in a lattice with the capability of hopping between sites and with on-site interactions, and is ruled by two key parameters: the hopping amplitude t (which allows particles to move between neighboring sites) and the on-site interaction strength U (which penalizes double occupancy). The competition between kinetic delocalization and interaction-induced localization gives rise to rich many-body phenomena, including *Mott* insulating phases and magnetism. [133]

In the limit of near unitary filling and strong interactions ($U \gg t$), doubly occupied sites are energetically suppressed, and the dynamics of the system can be effectively captured by projecting the Hubbard model into the subspace of no double occupancy. This projection leads to effective spin and density interactions between particles and results in the so-called " t - J - V model", which, in the isotropic case, includes:

- a hopping term (t) describing the kinetic energy of holes,
- a superexchange interaction term (J) arising from virtual tunneling processes and responsible for spin-spin correlations
- and optionally, a density-density interaction term (V) reflecting longer-range interactions beyond nearest-neighbor superexchange.

In several significant specific cases, such as spin-dependent tunneling, spin-orbit coupling, external fields, or constrained geometries, we will get XXZ-type models ruled by anisotropic terms, such as $t_{\uparrow}$, $t_{\downarrow}$, $J_{\perp}$ and J_z .

To better understand the role of such an interplay between anisotropic terms, a synthetic casistic can be envisaged:

- If $J_{\perp} = J_Z = J$: fully SU(2)-symmetric Heisenberg model.
- If $J_Z > J_{\perp}$: Ising-like (spin alignment favored along z).
- If $J_Z < J_{\perp}$: XY-like (spin alignment favored in the xy -plane).
- If $J_Z = 0$: pure XY model.
- If $J_{\perp} = 0$: pure Ising model.

Neglecting density-density interactions (V) leads to the simplified **t-J Model**, which is given as: [134]

$$\hat{H}_{tJ} = \hat{H}_t + \hat{H}_J, \quad (5.1)$$

$$\hat{H}_t = - \sum_{i < j} \sum_{\sigma=\downarrow, \uparrow} \frac{t_{\sigma}}{r_{ij}^3} \left(\hat{a}_{i,\sigma}^{\dagger} \hat{a}_{j,h}^{\dagger} \hat{a}_{i,h} \hat{a}_{j,\sigma} + \text{h.c.} \right), \quad (5.2)$$

$$\hat{H}_J = \sum_{i < j} \frac{1}{r_{ij}^6} \left[J_z \hat{S}_i^z \hat{S}_j^z + \frac{J_{\perp}}{2} \left(\hat{S}_i^+ \hat{S}_j^- + \text{h.c.} \right) \right], \quad (5.3)$$

where r_{ij} denotes the distance between particles i and j . Here, $\hat{a}_{i,\sigma}^{\dagger}$ creates a fermion (or boson) with spin $\sigma = \uparrow, \downarrow$ at site i , and $\hat{a}_{i,h}$ annihilates a hole at site i , defined as Gutzwiller-projected annihilation operator:

$$\hat{a}_{i,h}^{\dagger} = \sum_{\sigma} \hat{a}_{i,\sigma} (1 - \hat{n}_{i,\bar{\sigma}}), \quad (5.4)$$

with $\bar{\sigma}$ denoting the opposite spin and $\hat{n}_{i,\sigma} = \hat{a}_{i,\sigma}^{\dagger} \hat{a}_{i,\sigma}$ the number operator. This expression ensures that the hole operator acts only on singly occupied states, in accordance with the no-double-occupancy constraint.

The spin operators $\hat{S}_i^z$, $\hat{S}_i^+$, and $\hat{S}_i^-$ act within the subspace of singly occupied sites and describe the local spin degrees of freedom. They are defined in terms of the fermionic (or bosonic) creation and annihilation operators as:

$$\hat{S}_i^z = \frac{1}{2} (\hat{n}_{i,\uparrow} - \hat{n}_{i,\downarrow}), \quad \hat{S}_i^+ = \hat{a}_{i,\uparrow}^{\dagger} \hat{a}_{i,\downarrow}, \quad \hat{S}_i^- = \hat{a}_{i,\downarrow}^{\dagger} \hat{a}_{i,\uparrow}. \quad (5.5)$$

These operators satisfy the usual SU(2) spin algebra within the singly occupied subspace. In the projected Hilbert space (with no double occupancy), they effectively capture the exchange of spin states between neighboring particles and are responsible for magnon propagation in the system.

The t - J model thus contains two dominant **energy scales**: the hopping amplitude t , which rules the kinetic energy and propagation of holes, and the superexchange energy J , which arises from second-order tunneling processes and rules spin–spin interactions. This superexchange term J emerges as the effective low-energy theory of the Hubbard model in the large- U limit, with $J = \pm 4t^2/U$, where the sign depends on the particle statistics (positive for fermions, negative for bosons).




Figure 5.2: Phase diagram of the multi-particle bound states in the fermionic zigzag ladder for different J/t and different ratios of the magnon to hole densities. The vertical dashed line represents the point $J/t \sim 2$ at which the hole-magnon bound states disappears. The ∞ line represents the case of a single hole with an arbitrary number of magnons.

The phase diagram in Fig. 5.2 reveals the formation of several distinct multi-particle bound states, classified according to the number of holes (H) and magnons (M):

- **1H1M**: A mobile hole–magnon bound state. This is the primary and most robust composite excitation, stabilized purely by kinetic frustration.
- **1H2M**: A three-body bound state formed by one hole and two magnons, representing a many-body generalization of the kinetic binding mechanism.
- **2H3M**: An extended composite cluster with internal spin texture, characterized by nontrivial transverse and spiral-like spin correlations.

These objects behave as **emergent quasiparticles rather than static defects**: they are internally bound, while their center of mass gets delocalized and mobile across the lattice.

Geometry tuning and dimensional crossover

The studied zigzag ladder geometry consists of two legs with sites connected both along each leg and across the rungs in a staggered triangular pattern (Fig. 5.3). The lattice spacing between nearest neighbors on the same leg is denoted by a , while the vertical distance between the two legs is h . This forms a sequence of corner-sharing triangles, introducing geometric frustration into the system. The effective hopping amplitude along the two types of bonds (along-leg and diagonal inter-leg) depends on the relative orientation and distance between neighboring sites.




Figure 5.3: Study platform with 19 fermions prepared in a tunable zigzag ladder geometry. Thanks to a customizing a/h ratio we can high tune its the system dynamics. The dynamics of the system can be tuned through the a/h ratio.

In particular, there exists a *sweet spot* in the geometry where binding effects are strongly enhanced due to optimized interference of kinetic paths. At $h/a = 1/2$, destructive interference is maximized, leading to a significant increase - by nearly a factor of 15 - in the binding energy of hole-magnon pairs. However, this configuration ($h/a = 1/2$) also results in an overly strong tunneling amplitude, which can suppress localization and mask bound state signatures. A more favorable compromise is achieved at $h/a = 0.86(\sqrt{3}/2)$, where the binding energy is still enhanced by approximately a factor of 6, while maintaining balanced and symmetric tunneling amplitudes along both types of bonds. This choice offers an optimal regime for observing frustration-induced binding and serves as the geometric setting for our simulations (see Table 6).

h/a	a (μm)	Rung ($2\pi\times\text{MHz}$)				Leg ($2\pi\times\text{MHz}$)			
		$t_{\uparrow}$	$t_{\downarrow}$	$J_{\perp}$	J^z	$t_{\uparrow}$	$t_{\downarrow}$	$J_{\perp}$	J^z
0.86	14.7	-1.1	-1.2	0.107	-0.071	-0.95	-1.0	0.075	-0.05
0.5	19.6	-1.1	-1.2	0.110	-0.076	-0.40	-0.42	0.013	-0.008

Table 6: Tunable interaction parameters of the zigzag ladder geometry. By varying the ratio h/a and the lattice spacing a , independent control over rung and leg tunneling (t) and spin-exchange couplings ($J_{\perp}$, J^z) is achieved.

As a reference geometry, it was also investigated a fully two-dimensional triangular lattice comprising 37 sites arranged with uniform nearest-neighbor spacing. This configuration eliminates residual quasi-1D effects present in the ladder and represents the maximally frustrated case for kinetic motion. Measurements performed on this lattice serve as a benchmark for the ladder results, allowing us to disentangle genuine dimensional effects from finite-size and boundary-induced artifacts. In particular, the 2D lattice provides a stringent test of the robustness of hole-magnon binding, as well as a direct comparison of binding length scales, mobility, and correlation patterns across different dimensionalities.

The first analysis focuses on the kinetically induced bound state formed by a single hole and a single magnon (1H1M) on a triangular ladder geometry. In a fully spin-polarized bond, a hole can lower its kinetic energy by forming an antisymmetric wavefunction across the two sites, which maximizes its effective tunneling amplitude and thus its mobility [135]. The introduction of a magnon locally resolves this frustration. When the magnon forms an antisymmetric singlet state, $(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$, with a neighboring spin-down atom, it effectively reverses the sign of the hole's hopping amplitude, $t_{\text{eff}} = -t$. This sign change acts as an effective π -flux threading the plaquette. In the presence of this negative effective hopping, the hole minimizes its kinetic energy by adopting a symmetric wavefunction, which can be accommodated on the triangular plaquette without frustration.

As a consequence, the magnon lowers the kinetic energy of the hole, leading to the formation of a stable hole-magnon bound state. Importantly, this composite object remains highly mobile and can propagate through the lattice without suffering from kinetic frustration, as illustrated in Fig. 5.4.




Figure 5.4: Mechanism of kinetically-induced binding. (i) A magnon forms a singlet state with adjacent spins surrounding a hole, creating a bound state. (ii) The hole moves to magnon's position, an internal dynamic of the bound state that remains kinetically unfrustrated. (iii) The hole hops to a neighboring spin-down site, effectively moving the bound state by one lattice site

The binding strength can be quantified by the binding energy

$$E_b = E(1H1M) - E(1H) - E(1M) + E(0H0M), \quad (5.6)$$

where $E(nHmM)$ denotes the ground-state energy of a system containing n holes and m magnons. A negative value of E_b signals the existence of a bound state.

On triangular ladders, the binding energy depends sensitively on the lattice geometry. As shown in Fig. 5.5, tuning the ratio h/a reveals a pronounced minimum in the binding energy, reaching $E_b \simeq -1.5|t|$ at $h/a = 0.55$. As the system is continuously deformed toward a one-dimensional chain ($h/a \rightarrow 0$), where dipolar tunneling dominates, the binding weakens and the binding energy increases to $E_b \simeq -0.85|t|$.




Figure 5.5: Binding energy for a single hole in a fully polarized zigzag ladder as a function of the distance between the legs of that ladder. When the two legs are either overlapped or too far from each other we recover 1D chain conditions. For $h/a = 0.86$ the profile exhibits a flex and for $h/a = 0.55$ has the minimum. The two blue dots mark the two working distances: $h/a = 0.5, 0.86$

5.2 EXPERIMENTAL PREPARATION, EVOLUTION AND DETECTION

The discussed model is experimentally implemented by encoding the spin degrees of freedom and *holes* using Rydberg levels of ^{87}Rb atoms, with state mapping as follows:

- $|61S\rangle \equiv |\uparrow\rangle$ (spin-up = *Magnon*),
- $|60P\rangle \equiv |h\rangle$ (absence of excitation of spin = *Hole*),
- $|60S\rangle \equiv |\downarrow\rangle$ (spin-down).

These Rydberg levels are coherently coupled via microwave radiation, allowing for state preparation, spin manipulation, and control of excitations with site-resolved addressing when combined with optical potentials. The use of microwave-coupled Rydberg levels to define the effective spin-1/2 manifold allows us to treat the system within the $S^z = \pm 1/2$ picture.




Figure 5.6: (a) Mapping of a spin-1/2 system with a hole. Colors are assigned to label the states, following the convention used throughout the manuscript figures. (b) Measurements of final state occupations. The populations in the spin-down ($\hat{n}^\downarrow$), hole ($\hat{n}^h$), and spin-up ($\hat{n}^\uparrow$) bases are measured in separate experimental runs, where the state at each site is mapped to the presence or absence of an atom.

To resolve *hole-magnon* dynamics in real space, a fluorescence-based detection scheme has been employed leveraging state-dependent shelving and imaging of ground-state atoms in $|^5S_{1/2}\rangle$, as explained in Section 5.2.3.

5.2.1 Site-Resolved State Preparation

Single ^{87}Rb atoms are initially loaded from a $5\,\mu\text{K}$ magneto-optical trap (MOT) into optical tweezers, with a final temperature of $13\,\mu\text{K}$ [136]. Using a two-dimensional acousto-optic deflector (AOD), the atoms are rearranged into a defect-free array with the desired geometry. Following rearrangement, the atomic temperature is reduced through a combination of optical molasses and Raman sideband cooling. The atoms are then optically pumped into the stretched ground state $|g\rangle \equiv |5S_{1/2}, F=2, m_F=2\rangle$, using a σ^+ -polarized laser at $795\,\text{nm}$.

Prior to Rydberg excitation, the tweezer depth is adiabatically reduced by a factor of approximately 100 in order to minimize the momentum spread of the atomic wavefunctions. The optical tweezers are then switched off, and all atoms are coherently transferred to the Rydberg state $|\downarrow\rangle$ using a two-photon stimulated Raman adiabatic passage (STIRAP) involving lasers at $420\,\text{nm}$ and $1014\,\text{nm}$.

The preparation of the desired initial spin and hole configurations relies on a combination of global microwave pulses and site-resolved light shifts. Two off-resonant lasers are employed to generate differential light shifts between the relevant Rydberg states:

- A $1014\,\text{nm}$ laser red-detuned by $\Delta \simeq 2\pi \times 300\,\text{MHz}$ from the $|6P_{3/2}\rangle \leftrightarrow |60S_{1/2}\rangle$ transition produces a light shift $\delta_\downarrow \simeq 2\pi \times 25\,\text{MHz}$. This beam is patterned using a spatial light modulator (SLM), allowing for site-selective addressing across the array.
- A second laser, red-detuned by the same amount from the $|6P_{3/2}\rangle \leftrightarrow |61S_{1/2}\rangle$ transition, generates an additional light shift $\delta_\uparrow \simeq 2\pi \times 25\,\text{MHz}$ and is controlled using an acousto-optic modulator (AOM).

Based on the applied light shifts, atoms are classified into three groups:

- **Magnon sites** ($|\uparrow\rangle$): atoms experiencing only the $\delta_\downarrow$ shift;
- **Hole sites** ($|h\rangle$): atoms experiencing both $\delta_\downarrow$ and $\delta_\uparrow$ shifts;
- **Background sites** ($|\downarrow\rangle$): atoms experiencing no light shift.

The state initialization sequence begins with all atoms prepared in the $|\downarrow\rangle$ state. A global $20\,\text{ns}$ microwave π -pulse transfers the entire array to the hole state $|h\rangle$. Subsequently, the $\delta_\downarrow$ light shift is applied to the sites designated for magnons and holes. A $39\,\text{ns}$ microwave π -pulse, resonant with the $|h\rangle \leftrightarrow |\downarrow\rangle$ transition, transfers only the background




Figure 5.7: **(a)** The states used in the mapping of the t - J Hamiltonian are indicated in black. Each of the three columns represents one class of atoms: nonaddressed atoms (left column) are prepared in $|\downarrow\rangle$, atoms with both $\delta_\downarrow$ and $\delta_\uparrow$ light shifts (center column) are prepared in $|h\rangle$, and atoms with only the $\delta_\downarrow$ light shift (right column) are prepared in $|\uparrow\rangle$. **(b)** State Initialization and Quasi-Adiabatic Ramp. Following Rydberg excitation via STIRAP, CHADOQ implement a sequence of microwave pulses together with site-resolved light shifts, $\delta_\downarrow$ and $\delta_\uparrow$, to prepare the target initial product state. Subsequently, both light shifts are ramped down to adiabatically approach a low-energy state.

atoms back to $|\downarrow\rangle$, while the addressed atoms remain in $|h\rangle$ due to the induced detuning.

Finally, the $\delta_\uparrow$ light shift is applied to the designated hole sites, and a 75 ns microwave π -pulse resonant with the $|h\rangle \leftrightarrow |\uparrow\rangle$ transition is applied. This pulse selectively transfers atoms experiencing only $\delta_\downarrow$ to the $|\uparrow\rangle$ state, while atoms subjected to both light shifts remain in the hole state $|h\rangle$. This sequence prepares the desired product state with site-resolved control over spins and holes.

5.2.2 Quasi-Adiabatic Evolution to Low-Energy States

After preparation of the initial product state, the system is driven toward the low-energy manifold of the t - J Hamiltonian by quasi-adiabatically ramping down both light shifts. The ramp consists of two stages. First, the light shifts are rapidly reduced to approximately $2\pi \times 5\text{ MHz}$ within $1\text{ }\mu\text{s}$. At this stage, the remaining detuning is still sufficient to suppress dynamics driven by the bare Hamiltonian, ensuring adiabaticity.

Subsequently, both light shifts are exponentially decreased to zero with a time constant of $1\ \mu\text{s}$ using the AOMs.




Figure 5.8: Eigenstates overlap during the evolution. (a) The average distance between the hole and the magnon (blue) and the eigenenergy (grey) are shown as a function of the eigenstate index up to 25. We consider here the lowest 10 eigenstates as bound states, characterized by a short hole-magnon distance. (b) The time evolution of the overlap of the prepared state with the true ground state $|E_0\rangle$ (green) and with the subspace spanned by the bound states (red). While the final overlap with the ground state is approximately 60%, the overlap with the bound-state manifold is close to unity, confirming the successful preparation of the composite particle.

Due to the finite duration of the ramp, the evolution is not fully adiabatic with respect to the true many-body ground state. Numerical simulations of the experimental ramp profile indicate a final overlap of approximately 60% with the ground state (see Fig. 5.8). Nevertheless, the evolution is sufficiently slow to enter the low-energy manifold of bound states. Exact diagonalization reveals that the lowest ten eigenstates are characterized by an average hole-magnon separation of approximately one lattice site. Importantly, the overlap of the experimentally prepared state with the subspace spanned by these bound states is close to unity, demonstrating that the protocol efficiently prepares hole-magnon bound states with high fidelity, despite residual diabatic effects.

5.2.3 State Detection and Measurement Bases

State detection is performed by mapping selected internal states onto the presence or absence of atoms in the final fluorescence image acquired at the end of each experimental sequence (see Fig. 5.4(iii)).

Since the characteristic interaction timescale of the system is $2\pi/t \sim 1 \mu\text{s}$, it is essential to halt the many-body dynamics prior to readout in order to avoid any evolution during the measurement pulses.

Freezing of the Dynamics

To freeze the dynamics, two dedicated microwave pulses are applied to transfer atoms in the states $|h\rangle$ and $|\uparrow\rangle$ into auxiliary Rydberg manifolds, $|G\rangle \equiv |58G\rangle$, $|P'\rangle \equiv |61P\rangle$, which do not interact with atoms remaining in the $|\downarrow\rangle$ state (see Fig. 5.7a). This procedure isolates the populations to be measured and prevents further coherent dynamics during the detection stage.




Figure 5.9: **(a)** Sequence to measure the $|\downarrow\rangle$ population, $\hat{n}_{\downarrow}$. The process begins with “freezing” pulses that transfer the $|h\rangle$ and $|\uparrow\rangle$ states to the non-interacting $|58G\rangle$ and $|61P\rangle$ levels, respectively. **(b)** Similarly, to measure the population of $|\uparrow\rangle$ ($\hat{n}_{\uparrow}$) or $|h\rangle$ ($\hat{n}_h$), the sequence is modified to freeze the target state while de-exciting the other two Rydberg states, thereby mapping the target state onto the absence of an atom (see text). **(c)** The sequence for the one-hole–two-magnon (1H2M) experiment also measures $\hat{n}_{\uparrow}$, but is designed to map the $|\uparrow\rangle$ state to the presence of an atom.

Basis Measurement $|\downarrow\rangle$

Following the freezing pulses, a 1014 nm laser pulse resonant with the $|6P_{3/2}\rangle \leftrightarrow |\downarrow\rangle$ transition is applied for 6 μs . Atoms in the $|\downarrow\rangle$ state are deexcited through the short-lived $|6P_{3/2}\rangle$ state and decay into the ground-state manifold. The optical tweezers are then switched back on, recapturing ground-state atoms, while atoms remaining in Rydberg states are expelled via the ponderomotive force.

As a result, atoms that were initially in the $|\downarrow\rangle$ state are imaged, whereas atoms in $|h\rangle$ or $|\uparrow\rangle$ are lost. This procedure therefore realizes a projective measurement of the occupation operator $\hat{n}_{\downarrow}$. Since the fluorescence image provides a binary outcome (recaptured or lost), this measurement does not distinguish between $|h\rangle$ and $|\uparrow\rangle$; for this reason, we refer to it as a $|\downarrow\rangle$ -basis measurement. The protocol is illustrated in Fig. 5.9a.

Basis Measurements $|h\rangle$ and $|\uparrow\rangle$

To access the occupation of holes or magnons, the detection sequence is modified accordingly.

For the $|h\rangle$ -basis measurement (Fig. 5.9b), atoms in the $|\downarrow\rangle$ state are first deexcited to the ground state as described above. Subsequently, two microwave π -pulses followed by an additional deexcitation pulse transfer atoms initially in $|\uparrow\rangle$ to the ground-state manifold. As a result, atoms initially in $|\downarrow\rangle$ or $|\uparrow\rangle$ are recaptured and imaged, while atoms in the $|h\rangle$ state remain in a Rydberg manifold and are expelled. Interpreting lost atoms as holes allows direct measurement of the hole occupation operator $\hat{n}_h$.

Basis Measurements $|\uparrow\rangle$

The $|\uparrow\rangle$ -basis measurement (Fig. 5.9c) is realized by first deexciting atoms in $|\downarrow\rangle$, then transferring atoms in $|h\rangle$ to $|\downarrow\rangle$ via a microwave π -pulse, followed by a second 1014 nm deexcitation pulse. With this protocol, atoms initially in $|\downarrow\rangle$ or $|h\rangle$ are imaged, while atoms that were in $|\uparrow\rangle$ remain in Rydberg states and are lost. This effectively implements a measurement of the magnon occupation operator $\hat{n}_{\uparrow}$.

In both last two detection schemes, a final freezing pulse is applied to transfer atoms mapped onto absence in the final image into non-interacting Rydberg manifolds, ensuring that they remain expelled by the optical tweezers.

5.3 TRANSVERSE SPIN CORRELATION ANALYSIS

To investigate the impact of kinetic frustration and emergent magnetic order in doped ladders and lattices, we can analyze correlation functions that are sensitive to the relative motion and binding of holes and magnons. While transverse spin correlations probe coherent spin exchange in the XY plane, the most direct experimental signature of a kinetically induced bound state is provided by *connected density correlations* between the hole and the magnon.

A hallmark of binding between a hole and a magnon is a positive connected correlator

$$\langle \hat{n}_i^h \hat{n}_j^\uparrow \rangle_c = \langle \hat{n}_i^h \hat{n}_j^\uparrow \rangle - \langle \hat{n}_i^h \rangle \langle \hat{n}_j^\uparrow \rangle, \quad (5.7)$$

which isolates genuine correlations beyond statistical coincidences. A positive value indicates that the joint probability of finding a hole at site i and a magnon at site j exceeds that expected from uncorrelated occupations, signaling an effective attraction between the two constituents.

In our experimental implementation, single-site resolved projective detection does not allow us to distinguish all three local states $|\downarrow\rangle$, $|\uparrow\rangle$, and $|h\rangle$ in a single shot. As a consequence, direct access to the ordered correlator $\langle \hat{n}_i^h \hat{n}_j^\uparrow \rangle_c$ is not possible. However, exploiting permutation symmetry, we can statistically reconstruct a *symmetrized connected correlator*:

$$\langle \hat{n}_i^h \hat{n}_j^\uparrow \rangle_c^s = \frac{1}{2} \left(\langle \hat{n}_i^h \hat{n}_j^\uparrow \rangle_c + \langle \hat{n}_i^\uparrow \hat{n}_j^h \rangle_c \right). \quad (5.8)$$

This quantity faithfully captures hole-magnon correlations while remaining compatible with the experimental readout.




Figure 5.10: Time evolution of the hole-magnon connected correlations, $\langle \hat{n}_i^h \hat{n}_j^\uparrow \rangle_c^s$.

The symmetrized connected correlations are visualized as spatial correlation maps at different evolution times T , as shown in Figure 5.10. At $T = 0$, positive correlations reflect the initially prepared configuration, with the hole and magnon localized on neighboring sites. As the site-dependent light shifts $\delta_{\uparrow}$ and $\delta_{\downarrow}$ are ramped down, the pair delocalizes and correlations spread over the system.

Crucially, the persistence of positive correlations on the super- and sub-diagonals indicates that the hole and magnon predominantly occupy neighboring sites, while moving together as a composite object. This behavior is a direct manifestation of kinetic binding. At long times, the correlations remain localized around the center of the ladder, revealing that the final state consists of a mobile, but bound hole–magnon pair.

Center-of-mass distribution of the bound pair

To quantify the spatial extent and mobility of the composite object, we can compute the center-of-mass (COM) probability distribution as

$$\mathcal{P}_{\text{COM}}(d) = \sum_{(r_i+r_j)/2=d} \langle \hat{n}_i^h \hat{n}_j^{\uparrow} \rangle^s, \quad (5.9)$$

where the sum runs over all pairs of sites whose center of mass is located at position d . The measured COM distribution is well captured by that of a single particle confined within the ladder, approximately following a $\sin^2$ envelope, demonstrating that the bound hole–magnon pair behaves as an effective quasiparticle (Fig. 5.11).




Figure 5.11: Slice of the COM probability distribution, $\mathcal{P}_{\text{COM}}$, along the center of the ladder (see inset). Experimental data (blue circles) is compared to the corresponding $\sin^2$ distribution (solid grey line). The shaded area represents numerically simulated distribution including errors.

Further insight into the binding mechanism is obtained by averaging correlations over pairs of sites separated by the same distance d :

$$C_s(d) = \frac{1}{N_d} \sum_{|r_i - r_j| = d} \langle \hat{n}_i^h \hat{n}_j^\uparrow \rangle_s, \quad (5.10)$$

where N_d is the number of site pairs at separation d . The distance metric $|\cdot|$ depends on the geometry: site-index separation for the ladder and Manhattan distance for a two-dimensional lattice.

In both geometries, the correlations decay approximately exponentially with distance, reflecting the finite spatial extent of the bound state. This behavior can be interpreted as the evanescent tail of the bound-state wavefunction, analogous to the decay outside a finite potential well. The experimentally extracted decay lengths are larger than the theoretical ground-state values, a discrepancy attributed to diabaticity during the ramp and to state-preparation-and-measurement errors.




Figure 5.12: Top: Hole-magnon non-connected and symmetrized correlations $C_s(d)$ versus distance for ladders (blue circles) and 2D arrays (tan circles). Solid lines represent an exponential fit of the correlation. Shaded areas are the numerical simulation including errors. Bottom: Normalized connected correlations $C_{c,s}(d)/C_{c,s}(1)$. Dashed lines correspond to the theoretical ground state. The inset shows the dominant kinetic contributions of the binding, and the red plaquettes are kinetically frustrated while the green one's frustration is alleviated. Error bars denote one standard error estimated via bootstrap, and are smaller than marker size.

Figure 5.12 provides a detailed, distance-resolved characterization of the hole–magnon bound state in both the zigzag ladder and the two-dimensional triangular lattice.

The **top panel** shows the symmetrized, non-connected hole–magnon correlations $C_s(d)$ as a function of the relative distance d between the hole and the magnon. Blue symbols correspond to the ladder geometry and tan to 2D geometry. In both cases, the correlations decay approximately exponentially with distance, indicating that the hole and magnon remain spatially confined within a finite binding length. Solid lines show exponential fits to the experimental data, while the shaded regions represent numerical simulations that include experimental imperfections such as diabatic excitations and detection errors. The overall agreement confirms that the observed decay is a genuine signature of binding rather than a trivial consequence of finite system size or averaging.

The **bottom panel** displays the corresponding connected and symmetrized correlations, normalized to their nearest-neighbor value, namely $C_{c,s}(d)/C_{c,s}(1)$. By subtracting uncorrelated contributions, these observables isolate genuine correlations between the hole and the magnon. Dashed lines indicate the theoretical ground-state predictions. In the two-dimensional lattice, connected correlations become negative already at next-nearest-neighbor distances, signaling a strong spatial confinement of the bound pair. In contrast, in the ladder geometry, correlations remain positive over a larger range, reflecting a more extended bound state.

The inset schematically illustrates the microscopic origin of this dimensional effect. Red triangular plaquettes indicate regions where kinetic frustration suppresses hole motion, while light-green plaquette corresponds to configurations where frustration is relieved by the presence of a nearby magnon. In both geometries, nearest-neighbor hole–magnon pairing removes frustration on adjacent plaquettes. However, in two dimensions, separating the pair does not provide additional kinetic relief, leading to a stronger energetic penalty and a more compact bound state compared to the ladder.

Overall, Fig. 5.12 demonstrates that distance-resolved symmetrized and connected correlations constitute a powerful diagnostic of kinetically induced binding. When combined with center-of-mass distributions, these observables provide direct evidence for mobile composite quasiparticles and reveal how lattice geometry and dimensionality fundamentally shape their internal structure and spatial extent.

5.4 SUMMARY AND OUTLOOK

In this chapter, we investigated the manifestation of kinetically induced bound states in frustrated lattice geometries, focusing on the dynamics of holes and magnons described by an effective t - J Hamiltonian. Unlike conventional binding mechanisms driven by attractive interactions, the bound states studied here arise purely from kinetic frustration: destructive interference between multiple hopping paths suppresses single-particle motion, while correlated motion of composite objects restores kinetic energy and stabilizes bound states.

Using a tunable zigzag ladder geometry, CHADOQ demonstrated how geometric frustration can be continuously controlled through the ratio h/a , allowing access to regimes where hole-magnon binding is strongly enhanced. In particular, identifying an optimal compromise geometry at $h/a = \sqrt{3}/2$, where the binding energy remains large while preserving balanced tunneling amplitudes and favorable experimental conditions. A fully two-dimensional triangular lattice was also employed as a benchmark, providing a maximally frustrated reference system and highlighting genuine dimensional effects on the binding length and correlation structure.

Experimentally, the realization of the t - J model with microwave-coupled Rydberg states enabled site-resolved preparation, coherent evolution, and state-selective detection of holes and magnons. The observation of persistent, positive connected hole-magnon correlations, together with delocalized center-of-mass motion, provided unambiguous evidence for the formation of mobile composite quasiparticles. Distance-resolved correlation measurements further revealed the finite spatial extent of the bound states and their dependence on lattice geometry and dimensionality. Beyond the fundamental 1H1M bound state, the phase diagram revealed a hierarchy of multi-particle bound complexes involving multiple holes and magnons, pointing toward a broader tendency of frustrated systems to self-organize into mobile clusters. These results establish kinetically induced binding as a robust and experimentally accessible phenomenon and open new perspectives for exploring exotic many-body states driven by frustration rather than interactions.

Looking forward, the ability to engineer frustration, dimensionality, and spin anisotropy in Rydberg tweezer arrays offers a powerful platform for studying doped quantum magnets, emergent quasiparticles, and nontrivial transport phenomena. Extensions to larger systems, longer evolution times, and controlled disorder may further illuminate the role of kinetic frustration in strongly correlated quantum matter and its connections to phenomena such as unconventional superconductivity and spin-liquid behavior. CHADOQ's work shows that kinetic frustration alone can generate stable multi-particle bound states (1H1M, 1H2M) in a zigzag ladder, observable as mobile composite quasiparticles with finite center-of-mass dispersion.

CONCLUSIONS AND FUTURE PERSPECTIVES

In this thesis, I presented the design, development, and characterization of a neutral-atom quantum simulation platform based on individually trapped strontium atoms in optical tweezer arrays. The experimental apparatus was conceived, built, and optimized entirely within our laboratory, encompassing the construction of the ultra-high-vacuum system, optical setups, and control electronics, as well as the development of cooling, trapping, and imaging techniques. The overarching goal of this effort was to establish a robust, scalable, and versatile platform for the deterministic preparation and coherent control of single atoms for quantum simulation.

A key achievement of this work was the realization and optimization of a two-stage, three-dimensional magneto-optical trap (MOT) for ^{88}Sr . In the first stage, a “blue MOT” operating on the strong dipole-allowed $^1\text{S}_0 \rightarrow ^1\text{P}_1$ transition at 461 nm was used to capture and pre-cool atoms from a thermal beam, reaching temperatures of 6.4 mK and densities of $\sim 4 \times 10^9 \text{ cm}^{-3}$. By implementing a “red-shielded” blue MOT configuration through the addition of a beam resonant with the 689 nm transition, we enhanced the atom number by up to 80%. Subsequent transfer to a “red MOT” operating on the narrow intercombination line $^1\text{S}_0 \rightarrow ^3\text{P}_1$ at 689 nm enabled further cooling. Using a sequence combining broadband and single-frequency narrow-line cooling, we achieved final temperatures of 5.4 μK and densities approaching $3 \times 10^{10} \text{ cm}^{-3}$, providing excellent initial conditions for loading atoms into tightly focused optical tweezers.

Building on this cold atomic reservoir, we demonstrated efficient loading of atoms into optical tweezers and high-fidelity detection of single trapped atoms. Fluorescence imaging based on photon-counting detection yielded a single-atom detection fidelity of 99.986% and a post-imaging survival probability of approximately 97(2)%. Independent release-and-recapture measurements indicated atomic temperatures of 12.95(5) μK inside the tweezers, confirming that the cooling and loading protocols

preserve precise control over the atomic motional degrees of freedom. Furthermore, the observed vacuum-limited lifetime exceeding seven minutes is on par with, or better than, similar optical tweezer systems [6, 72], highlighting the quality of the ultra-high-vacuum environment achieved in our setup.

Together, these results constitute an excellent foundation for the realization of a fully programmable quantum simulator based on neutral strontium atoms arranged in reconfigurable optical tweezer arrays. The platform combines the geometric flexibility intrinsic to optical tweezers with the rich internal structure of two-electron atoms. Crucially, by exploiting coherent excitation to Rydberg states, the system enables the engineering of controllable, long-range interactions between individual atoms, opening access to a wide class of interacting many-body Hamiltonians.

SHORT-TERM GOALS: REORDERING AND RYDBERG INTERACTIONS

In the near future, a key upgrade will be the integration of a spatial light modulator (SLM), currently under characterization, for the generation of static tweezer arrays. Two crossed acousto-optic deflectors (AODs) will be used to dynamically transport individual atoms, enabling active reordering and the deterministic assembly of defect-free atomic arrays with arbitrary geometries. In parallel, we are exploring the use of a fast-refresh-rate SLM (~ 1.2 kHz) to implement parallel rearrangement protocols in two- and potentially three-dimensional configurations.

A major next step will be the implementation of a two-photon excitation scheme to access high-lying Rydberg states. This capability will introduce tunable, long-range interactions between atoms, marking the transition from passive trapping and control to the realization of fully interacting quantum systems. By tuning the interatomic distance, the principal quantum number, and the orbital character of the Rydberg states, it will be possible to engineer interactions ranging from isotropic van der Waals to strongly anisotropic dipole–dipole couplings. This high degree of tunability provides a powerful toolbox for engineering diverse quantum Hamiltonians and exploring a broad range of many-body phenomena.

QUANTUM SIMULATION OF SPIN MODELS

One of the most promising research directions enabled by this platform is the study of quantum spin systems. Optical tweezer arrays of Rydberg atoms are ideally suited to simulate Ising, XY, and more general spin Hamiltonians in one, two, and three dimensions. The mapping between atomic ground and Rydberg states and effective spin-1/2 degrees of freedom allows direct implementation of transverse-field Ising models with long-range interactions, as well as their extensions via spatially dependent driving fields and tailored geometries.

Beyond the Ising model, resonant dipole–dipole interactions between different Rydberg states enable the realization of XY-type Hamiltonians with coherent spin-exchange dynamics. These systems offer access to rich physical regimes, including quantum phase transitions, frustrated magnetism, symmetry-protected topological phases, and exotic quantum liquids. The single-site resolution and full control available in the present platform make it particularly well suited for studying non-equilibrium dynamics, disorder effects, and the interplay between geometry and interactions.

A NEW FRONTIER: KINETICALLY FRUSTRATED DYNAMICS

The work presented in the final chapter of this thesis places this experimental platform within a broader and emerging frontier of quantum simulation: the study of *kinetically frustrated* many-body systems. By exploiting frustrated lattice geometries and controlled doping with holes and spin excitations, we demonstrated that kinetic frustration alone can give rise to robust, mobile multi-particle bound states, such as hole–magnon composites, even in the absence of attractive interactions.

These results establish optical tweezer arrays of Rydberg atoms as a powerful platform for simulating doped quantum magnets and frustration-driven phenomena. The ability to engineer lattice geometry, dimensionality, spin anisotropy, and interaction range with single-particle resolution opens new pathways for investigating emergent quasiparticles, composite excitations, and nontrivial transport phenomena that are difficult to access with classical computational methods or traditional condensed-matter systems.

OUTLOOK

Although the primary focus of this thesis has been the development of a neutral-atom tweezer platform for quantum simulation, the same architecture is highly relevant for other frontier applications. In particular, it provides a promising route toward scalable quantum computing based on Rydberg-mediated entangling gates, as well as for quantum metrology and optical clock experiments with alkaline-earth atoms.

Overall, the results reported in this thesis demonstrate the successful implementation of a robust, high-fidelity platform for trapping, imaging, and controlling single strontium atoms. By combining precise single-particle control with scalable architectures and tunable long-range interactions, this platform lays the foundation for a new generation of quantum simulators capable of addressing open questions in strongly correlated quantum matter.

ACKNOWLEDGEMENTS

Once again, I find myself completing an important stage of my life, and, as it should be, I feel the urge to give thought to the people who made this journey possible. I have always sought to express my gratitude and affection toward those who have been kind to me, supportive, inspiring, and a source of dialogue and thoughtful discussion. However, words are not always enough and they rarely endure, and I believe this is precisely why the most heartfelt thanks should be put into writing.

First of all, I would like to thank the entire Department of Physics and the LENS Institute for creating a functional and welcoming working environment, in which everyone can find their own space and contribute.

I am deeply grateful to my mentors, Vladislav and Leonardo, for the patience, kindness, and wisdom with which they guided me through these three long years. I am sincerely thankful to my fellow travelers, Luca and Veronica, who made every day in the laboratory an adventure filled with laughter, challenges, and discoveries. I also thank my colleagues Jacopo, Andreas, Shawn, Giacomo, and the entire Quantum Gas Group.

In the end, as in the best stories, beautiful things are truly meaningful only when they are shared with special people. I thank my “physicist” friends, my “metaphysician” friends, my companions at Nuovo Teatro Tempesta, and above all my family, whose gentle presence alone allowed me to catch my breath whenever I needed it.

BIBLIOGRAPHY

- [1] Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S. Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682):579–584, 11 2017.
- [2] Antoine Browaeys and Thierry Lahaye. Many-body physics with individually controlled rydberg atoms. *Nature Physics*, 16(2):132–142, 2020.
- [3] Heng Shen and Jing Zhang. Entanglement-enhanced quantum metrology with neutral atom arrays. *National Science Review*, page nwaf149, 2025.
- [4] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with rydberg atoms. *Reviews of Modern Physics*, 82(3):2313–2363, Aug 2010.
- [5] Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Tout T. Wang, Sepehr Ebadi, Hannes Bernien, Markus Greiner, Vladan Vuletić, Hannes Pichler, and Mikhail D. Lukin. Parallel implementation of high-fidelity multiqubit gates with neutral atoms. *Physical Review Letters*, 123:170503, Oct 2019.
- [6] Ivaylo S. Madjarov, Jacob P. Covey, Adam L. Shaw, Joonhee Choi, Anant Kale, Alexandre Cooper, Hannes Pichler, Vladimir Schkolnik, Jason R. Williams, and Manuel Endres. High-fidelity entanglement and detection of alkaline-earth rydberg atoms. *Nature Physics*, 16(8):857–861, 2020.
- [7] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, M. Kwon, M. Ebert, J. Cherek, M. T. Lichtman, M. Gillette, J. Gilbert, D. Bowman, T. Ballance, C. Campbell, E. D. Dahl, O. Crawford, N. S. Blunt, B. Rogers, T. Noel, and M. Saffman. Multi-qubit entanglement and algorithms on a neutral-atom quantum computer. *Nature*, 604(7906):457–462, 2022.

- [8] Iulia M Georgescu, Sahel Ashhab, and Franco Nori. Quantum simulation. *Reviews of Modern Physics*, 86(1):153–185, 2014.
- [9] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, 1982.
- [10] S. Lloyd. Universal quantum simulators. *Science*, 273:1073–1078, 1996.
- [11] M. Vyalı A. Kitaev, A. Shen. Classical and quantum computation. 2002.
- [12] J. Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.
- [13] Richard P Feynman. Simulating physics with computers. In *Feynman and computation*, pages 133–153. cRc Press, 2018.
- [14] Iulia Buluta and Franco Nori. Quantum simulators. *Science*, 326(5949):108–111, 2009.
- [15] Ehud Altman, Kenneth R Brown, Giuseppe Carleo, Lincoln D Carr, Eugene Demler, Cheng Chin, Brian DeMarco, Sophia E Economou, Mark A Eriksson, Kai-Mei C Fu, et al. Quantum simulators: Architectures and opportunities. *PRX quantum*, 2(1):017003, 2021.
- [16] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.
- [17] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W Hänsch, and Immanuel Bloch. Quantum phase transition from a superfluid to a mott insulator in a gas of ultracold atoms. *nature*, 415(6867):39–44, 2002.
- [18] Anton Mazurenko, Christie S Chiu, Geoffrey Ji, Maxwell F Parsons, Márton Kanász-Nagy, Richard Schmidt, Fabian Grusdt, Eugene Demler, Daniel Greif, and Markus Greiner. A cold-atom fermi-hubbard antiferromagnet. *Nature*, 545(7655):462–466, 2017.
- [19] Jonathan Simon, Waseem S Bakr, Ruichao Ma, M Eric Tai, Philipp M Preiss, and Markus Greiner. Quantum simulation of antiferromagnetic spin chains in an optical lattice. *Nature*, 472(7343):307–312, 2011.

- [20] Gregor Jotzu, Michael Messer, Rémi Desbuquois, Martin Lebrat, Thomas Uehlinger, Daniel Greif, and Tilman Esslinger. Experimental realization of the topological haldane model with ultracold fermions. *Nature*, 515(7526):237–240, 2014.
- [21] Marc Cheneau, Peter Barmettler, Dario Poletti, Manuel Endres, Peter Schauß, Takeshi Fukuhara, Christian Gross, Immanuel Bloch, Corinna Kollath, and Stefan Kuhr. Light-cone-like spreading of correlations in a quantum many-body system. *Nature*, 481(7382):484–487, 2012.
- [22] Stefan Trotzky, Yu-Ao Chen, Andreas Flesch, Ian P McCulloch, Ulrich Schollwöck, Jens Eisert, and Immanuel Bloch. Probing the relaxation towards equilibrium in an isolated strongly correlated one-dimensional bose gas. *Nature physics*, 8(4):325–330, 2012.
- [23] Hartmut Häffner, Christian F Roos, and Rainer Blatt. Quantum computing with trapped ions. *Physics reports*, 469(4):155–203, 2008.
- [24] Christopher J Ballance, Thomas P Harty, Nobert M Linke, Martin A Sepiol, and David M Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Physical review letters*, 117(6):060504, 2016.
- [25] Diego Porras and J Ignacio Cirac. Effective quantum spin systems with trapped ions. *Physical review letters*, 92(20):207901, 2004.
- [26] Axel Friedenauer, Hector Schmitz, Jan Tibor Glueckert, Diego Porras, and Tobias Schätz. Simulating a quantum magnet with trapped ions. *Nature Physics*, 4(10):757–761, 2008.
- [27] Kihwan Kim, M-S Chang, Simcha Korenblit, Rajibul Islam, Emily E Edwards, James K Freericks, G-D Lin, L-M Duan, and Christopher Monroe. Quantum simulation of frustrated ising spins with trapped ions. *Nature*, 465(7298):590–593, 2010.
- [28] W. Morong, K.S. Collins, A. De, E. Stavropoulos, T. You, and C. Monroe. Engineering dynamically decoupled quantum simulations with trapped ions. *PRX Quantum*, 4:010334, Mar 2023.
- [29] Christopher Monroe. Quantum simulation of frustrated ising spins with trapped ions. In *APS March Meeting Abstracts*, volume 2010, pages X2–005, 2010.

- [30] Morten Kjaergaard, Mollie E Schwartz, Jochen Braumüller, Philip Krantz, Joel I-J Wang, Simon Gustavsson, and William D Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11(1):369–395, 2020.
- [31] Andrew A Houck, Hakan E Türeci, and Jens Koch. On-chip quantum simulation with superconducting circuits. *Nature Physics*, 8(4):292–299, 2012.
- [32] Alexandre Blais, Arne L Grimsmo, Steven M Girvin, and Andreas Wallraff. Circuit quantum electrodynamics. *Reviews of Modern Physics*, 93(2):025005, 2021.
- [33] David Marcos, Peter Rabl, Enrique Rico, and Peter Zoller. Superconducting circuits for quantum simulation of dynamical gauge fields. *arXiv preprint arXiv:1306.1674*, 2013.
- [34] Yves Salathé, Mintu Mondal, Markus Oppliger, Johannes Heinsoo, Philipp Kurpiers, Anton Potočnik, Antonio Mezzacapo, Urtzi Las Heras, Lucas Lamata, Enrique Solano, et al. Digital quantum simulation of spin models with circuit quantum electrodynamics. *Physical Review X*, 5(2):021027, 2015.
- [35] Alan Aspuru-Guzik and Philip Walther. Photonic quantum simulators. *Nature Physics*, 8(4):285–291, Apr 2012.
- [36] Jianwei Wang, Fabio Sciarrino, Anthony Laing, and Mark G. Thompson. Integrated photonic quantum technologies. *Nature Photonics*, 2019.
- [37] Daniel Loss and David P DiVincenzo. Quantum computation with quantum dots. *Physical Review A*, 57(1):120, 1998.
- [38] Ronald Hanson, Leo P Kouwenhoven, Jason R Petta, Seigo Tarucha, and Lieven MK Vandersypen. Spins in few-electron quantum dots. *Reviews of modern physics*, 79(4):1217–1265, 2007.
- [39] Adam M Kaufman and Kang-Kuen Ni. Quantum science with optical tweezer arrays of ultracold atoms and molecules. *Nature Physics*, 17(12):1324–1333, 2021.
- [40] Manuel Endres, Hannes Bernien, Alexander Keesling, Harry Levine, Eric R. Anschuetz, Alexandre Krajenbrink, Crystal Senko, Vladan Vuletić, Markus Greiner, and Mikhail D. Lukin. Atom-by-atom

- assembly of defect-free one-dimensional cold atom arrays. *Science*, 354(6315):1024–1027, 2016.
- [41] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin. Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572):1242–1247, 2021.
- [42] Sepehr Ebadi, Tout T. Wang, Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Dolev Bluvstein, Rhine Samajdar, Hannes Pichler, Wen Wei Ho, Soonwon Choi, Subir Sachdev, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595(7866):227–232, 2021.
- [43] M. A. Norcia, A. W. Young, and A. M. Kaufman. Microscopic control and detection of ultracold strontium in optical-tweezer arrays. *Physical Review X*, 8:041054, Dec 2018.
- [44] Daichi Okuno, Yuma Nakamura, Toshi Kusano, Yosuke Takasu, Nobuyuki Takei, Hideki Konishi, and Yoshiro Takahashi. High-resolution spectroscopy and single-photon rydberg excitation of reconfigurable ytterbium atom tweezer arrays utilizing a metastable state. *Journal of the Physical Society of Japan*, 91(8):084301, 2022.
- [45] Pascal Scholl, Michael Schuler, Hannah J. Williams, Alexander A. Eberharter, Daniel Barredo, Kai-Niklas Schymik, Vincent Lienhard, Louis-Paul Henry, Thomas C. Lang, Thierry Lahaye, Andreas M. Läuchli, and Antoine Browaeys. Quantum simulation of 2d antiferromagnets with hundreds of rydberg atoms. *Nature*, 595(7866):233–238, 2021.
- [46] Grégoire Pichard, Desiree Lim, Étienne Bloch, Julien Vaneecloo, Lilian Bourachot, Gert-Jan Both, Guillaume Mériaux, Sylvain Dutartre, Richard Hostein, Julien Paris, Bruno Ximenez, Adrien Signoles, Antoine Browaeys, Thierry Lahaye, and Davide Dreon. Rearrangement of individual atoms in a 2000-site optical-tweezer array at cryogenic temperatures. *Physical Review Applied*, 22:024073, Aug 2024.
- [47] H.J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer, New York, 1999.

- [48] Claude Cohen-Tannoudji, Jacques Dupont-Roc, and Gilbert Grynberg. *Atom-Photon Interactions*. Wiley-VCH, Weinheim, 1998.
- [49] Arthur Ashkin. Trapping of atoms by resonance radiation pressure. *Physical Review Letters*, 40(12):729, 1978.
- [50] Stanley H Autler and Charles H Townes. Stark effect in rapidly varying fields. *Physical Review*, 100(2):703, 1955.
- [51] Rudolf Grimm, Matthias Weidemüller, and Yurii B. Ovchinnikov. Optical dipole traps for neutral atoms. In Benjamin Bederson and Herbert Walther, editors, *Advances in Atomic, Molecular, and Optical Physics*, volume 42 of *Advances In Atomic, Molecular, and Optical Physics*, pages 95–170. Academic Press, 2000.
- [52] Daniel A. Steck. Quantum and atom optics, October 2007. revision 0.11.3, 9 October 2016.
- [53] Martin M Boyd. *High precision spectroscopy of strontium in an optical lattice: Towards a new standard for frequency and time*. PhD thesis, University of Colorado at Boulder, 2007.
- [54] Ivan H Deutsch and Poul S Jessen. Quantum-state control in optical lattices. *Physical Review A*, 57(3):1972, 1998.
- [55] Alexander Urech, Ivo HA Knottnerus, Robert JC Spreeuw, and Florian Schreck. Narrow-line imaging of single strontium atoms in shallow optical tweezers. *Physical Review Research*, 4(2):023245, 2022.
- [56] Bahaa EA Saleh and Malvin Carl Teich. *Fundamentals of photonics, 2 volume set*. John Wiley & sons, 2019.
- [57] Susanne Quabis, Ralf Dorn, M Eberler, O Glöckl, and Gerd Leuchs. Focusing light to a tighter spot. *Optics communications*, 179(1-6):1–7, 2000.
- [58] Max Born and Emil Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.
- [59] Nicolas Schlosser, Georges Reymond, Igor Protsenko, and Philippe Grangier. Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature*, 411(6841):1024–1027, 2001.

- [60] Ivaylo Sashkov Madjarov. *Entangling, controlling, and detecting individual strontium atoms in optical tweezer arrays*. California Institute of Technology, 2021.
- [61] Jacob F. Sherson, Christof Weitenberg, Manuel Endres, Marc Cheneau, Immanuel Bloch, and Stefan Kuhr. Single-atom-resolved fluorescence imaging of an atomic mott insulator. *Nature*, 467(7311):68–72, Sep 2010.
- [62] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner. Probing the superfluid-to-mott insulator transition at the single-atom level. *Science*, 329(5991):547–550, 2010.
- [63] Marshall T DePue, S Lukman Winoto, D.J Han, and David S Weiss. Transient compression of a {MOT} and high intensity fluorescent imaging of optically thick clouds of atoms. *Optics Communications*, 180(1–3):73–79, 2000.
- [64] M. A. Norcia, H. Kim, W. B. Cairncross, M. Stone, A. Ryou, M. Jaffe, M. O. Brown, K. Barnes, P. Battaglino, A. Brown, K. Cassella, C. A. Chen, R. Coxe, D. Crow, J. Epstein, C. Griger, E. Halperin, F. Hummel, A. M. W. Jones, J. M. Kindem, J. King, K. Kotru, J. Lauigan, M. Li, M. Lu, E. Megidish, J. Marjanovic, M. McDonald, T. Mitiga, J. A. Muniz, S. Narayanaswami, C. Nishiguchi, T. Paule, K. A. Pawlak, L. S. Peng, K. L. Pudenz, A. Smull, D. Stack, M. Urbanek, R. J. M. van de Veerdonk, Z. Vendeiro, L. Wadleigh, T. Wilkason, T. Y. Wu, X. Xie, E. Zaly-Geller, X. Zhang, and B. J. Bloom. Iterative assembly of ^{171}Yb atom arrays in cavity-enhanced optical lattices. *arXiv e-prints*, page arXiv:2401.16177, Jan 2024.
- [65] Giovanna Morigi, Jürgen Eschner, JI Cirac, and P Zoller. Laser cooling of two trapped ions: Sideband cooling beyond the lamb-dicke limit. *Physical Review A*, 59(5):3797, 1999.
- [66] John D Teufel, Tobias Donner, Dale Li, Jennifer W Harlow, MS Allman, Katarina Cicak, Adam J Sirois, Jed D Whittaker, Konrad W Lehnert, and Raymond W Simmonds. Sideband cooling of micromechanical motion to the quantum ground state. *Nature*, 475(7356):359–363, 2011.
- [67] A. M. Kaufman, B. J. Lester, and C. A. Regal. Cooling a single atom in an optical tweezer to its quantum ground state. *Physical Review X*, 2:041014, Nov 2012.

- [68] Jacob P. Covey, Ivaylo S. Madjarov, Alexandre Cooper, and Manuel Endres. 2000-times repeated imaging of strontium atoms in clock-magic tweezer arrays. *Physical Review Letters*, 122:173201, May 2019.
- [69] Renhao Tao, Maximilian Ammenwerth, Flavien Gyger, Immanuel Bloch, and Johannes Zeiher. High-fidelity detection of large-scale atom arrays in an optical lattice. *Physical Review Letters*, 133:013401, Jul 2024.
- [70] Chun-Chia Chen, Shayne Bennetts, Rodrigo González Escudero, Florian Schreck, and Benjamin Pasquiou. Sisyphus optical lattice decelerator. *arXiv preprint arXiv:1810.07157*, 2018.
- [71] Vladyslav V Ivanov and Subhadeep Gupta. Laser-driven sisyphus cooling in an optical dipole trap. *Physical Review A—Atomic, Molecular, and Optical Physics*, 84(6):063417, 2011.
- [72] Alec Jenkins, Joanna W. Lis, Aruku Senoo, William F. McGrew, and Adam M. Kaufman. Ytterbium nuclear-spin qubits in an optical tweezer array. *Physical Review X*, 12:021027, May 2022.
- [73] Simon Stellmer, Florian Schreck, and Thomas C. Killian. *Degenerate Quantum Gases Of Strontium*, chapter CHAPTER 1, pages 1–80. 2014.
- [74]
- [75] Masami Yasuda, Tetsuo Kishimoto, Masao Takamoto, and Hidetoshi Katori. Photoassociation spectroscopy of sr 88: Reconstruction of the wave function near the last node. *Physical Review A—Atomic, Molecular, and Optical Physics*, 73(1):011403, 2006.
- [76] J. E. Sansonetti and G. Nave. Wavelengths, transition probabilities, and energy levels for the spectrum of neutral strontium (sri). *Journal of Physical and Chemical Reference Data*, 39(3):033103, 2010.
- [77] Matthew A Norcia, Julia RK Cline, Juan A Muniz, John M Robinson, Ross B Hutson, Akihisa Goban, G Edward Marti, Jun Ye, and James K Thompson. Frequency measurements of superradiance from the strontium clock transition. *Physical Review X*, 8(2):021036, 2018.
- [78] Alexey V Taichenachev, VI Yudin, CW Oates, CW Hoyt, ZW Barber, and L Hollberg. Magnetic field-induced spectroscopy of forbidden

optical transitions with application to lattice-based optical atomic clocks. *Physical review letters*, 96(8):083001, 2006.

- [79] Thomas F. Gallagher. *Rydberg Atoms*. Cambridge University Press, 1994.
- [80] Jongseok Lim, Han-gyeol Lee, and Jaewook Ahn. Review of cold rydberg atoms and their applications. *Journal of the Korean Physical Society*, 63(4):867–876, Aug 2013.
- [81] N. Šibalić, J.D. Pritchard, C.S. Adams, and K.J. Weatherill. Arc: An open-source library for calculating properties of alkali rydberg atoms. *Computer Physics Communications*, 220:319–331, 2017.
- [82] T. F. Gallagher and W. E. Cooke. Interactions of blackbody radiation with atoms. *Phys. Rev. A*, 21:588, 1980.
- [83] J. W. Farley and W. H. Wing. Accurate calculation of dynamic stark shifts and depopulation rates of rydberg energy levels induced by blackbody radiation. *Phys. Rev. A*, 23:2397, 1981.
- [84] I. I. Beterov, D. B. Tretyakov, I. I. Ryabtsev, V. M. Entin, A. Ekers, and N. N. Bezuglov. Ionization of rydberg atoms by blackbody radiation. *New Journal of Physics*, 11:013052, 2009.
- [85] L. Béguin, A. Vernier, R. Chicireanu, T. Lahaye, and A. Browaeys. Direct measurement of the van der waals interaction between two rydberg atoms. *Physical Review Letters*, 110:263201, Jun 2013.
- [86] Harry Levine, Alexander Keesling, Ahmed Omran, Hannes Bernien, Sylvain Schwartz, Alexander S. Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. High-fidelity control and entanglement of rydberg-atom qubits. *Physical Review Letters*, 121:123603, Sep 2018.
- [87] I Mourachko, D Comparat, Ferdinando de Tomasi, A Fioretti, P Nosbaum, VM Akulin, and P Pillet. Many-body effects in a frozen rydberg gas. *Physical Review Letters*, 80(2):253, 1998.
- [88] II Ryabtsev, DB Tretyakov, II Beterov, and VM Entin. Observation of the stark-tuned förster resonance between two rydberg atoms. *Physical review letters*, 104(7):073003, 2010.

- [89] II Beterov and M Saffman. Rydberg blockade, förster resonances, and quantum state measurements with different atomic species. *Physical Review A*, 92(4):042710, 2015.
- [90] Alpha Gaëtan, Yevhen Miroshnychenko, Tatjana Wilk, Amodsen Chotia, Matthieu Viteau, Daniel Comparat, Pierre Pillet, Antoine Browaeys, and Philippe Grangier. Observation of collective excitation of two individual atoms in the rydberg blockade regime. *Nature Physics*, 5(2):115–118, 2009.
- [91] Yevhen Miroshnychenko, A Gaëtan, C Evellin, P Grangier, D Comparat, P Pillet, T Wilk, and A Browaeys. Coherent excitation of a single atom to a rydberg state. *Physical Review A—Atomic, Molecular, and Optical Physics*, 82(1):013405, 2010.
- [92] Chang Qiao, Can-Zhu Tan, Julia Siegl, Fa-Chao Hu, Zhi-Jing Niu, Y. H. Jiang, Matthias Weidemüller, and Bing Zhu. Rydberg blockade in an ultracold strontium gas revealed by two-photon excitation dynamics. *arXiv e-prints*, page arXiv:2103.04528, Mar 2021.
- [93] Chenlong Zhang, Nicholas C. Anzalone, Rodrigo P. Faria, and Joshua M. Pearce. Open-source 3d-printable optics equipment. *PLoS ONE*, 8(3):1–13, 03 2013.
- [94] Luc Couturier, Ingo Nosske, Fachao Hu, Canzhu Tan, Chang Qiao, Y. H. Jiang, Peng Chen, and Matthias Weidemüller. Measurement of the strontium triplet rydberg series by depletion spectroscopy of ultracold atoms. *Physical Review A*, 99:022503, Feb 2019.
- [95] R. Ding, J. D. Whalen, S. K. Kanungo, T. C. Killian, F. B. Dunning, S. Yoshida, and J. Burgdörfer. Spectroscopy of ^{87}Sr triplet rydberg states. *Physical Review A*, 98:042505, Oct 2018.
- [96] Veronica Giardini, Luca Guariento, Andrea Fantini, Shawn Storm, Massimo Inguscio, Jacopo Catani, Giacomo Cappellini, Vladislav Gavryusev, and Leonardo Fallani. Single sr atoms in optical tweezer arrays for quantum simulation. *Atoms*, 14(1):1, 2025.
- [97] Luca Guariento. *A new setup for single Strontium atoms in Optical Tweezers*. PhD thesis, PhD degree in Quantum Technologies, 2025.
- [98] Veronica Giardini. *Development of a novel platform for quantum simulation with Sr Rydberg atoms*. PhD thesis, November 2025.

- [99] T. Denecker, Y. T. Chew, O. Guillemant, G. Watanabe, T. Tomita, K. Ohmori, and S. de Léséleuc. Measurement and feedforward correction of the fast phase noise of lasers. *Physical Review A*, 111:042614, Apr 2025.
- [100] Gianni Di Domenico, Stéphane Schilt, and Pierre Thomann. Simple approach to the relation between laser frequency noise and laser line shape. *Applied optics*, 49(25):4801–4807, 2010.
- [101] Theodor W Haensch. Tunable lasers and coherent light techniques for high resolution ultraviolet spectroscopy. Technical report, 1977.
- [102] Jitendra Nath Roy and Ashif Raja. Acousto-optic modulator based multi-operational logic circuit: design and analysis. *Journal of Optics*, pages 1–11, 2024.
- [103] Akis P Goutzoulis. *Design and fabrication of acousto-optic devices*. CRC Press, 2021.
- [104] Joseph W Goodman. *Introduction to Fourier optics*. Roberts and Company publishers, 2005.
- [105] B Zhu, J Schachenmayer, M Xu, F Herrera, J G Restrepo, M J Holland, and A M Rey. Synchronization of interacting quantum dipoles. *New Journal of Physics*, 17(8):083063, 2015.
- [106] Wei-Wei Liu, Yu-Xuan Ren, Hong-Fang Gao, Qing Sun, Zi-Qiang Wang, and Yin-Mei Li. Aberrations in holographic array optical tweezers corrected with zernike polynomials. *Acta Physica Sinica*, 61(18):188701, 2012.
- [107] Yi He, Mingdi Bao, Yiwei Chen, Hong Ye, Jinyu Fan, and Guohua Shi. Accuracy characterization of shack–hartmann sensor with residual error removal in spherical wavefront calibration. *Light: Advanced Manufacturing*, 4(4):393–403, 2024.
- [108] R. W. Gerchberg and W. O. Saxton. A practical algorithm for the determination of plane from image and diffraction pictures. *Optik*, 35(2):237–246, 1972. CiNii Research ID: 1572261550522209664.
- [109] Simon P Poland, Nikola Krstajić, Robert D Knight, Robert K Henderson, and Simon M Ameer-Beg. Development of a doubly weighted gerchberg–saxton algorithm for use in multibeam imaging applications. *Optics letters*, 39(8):2431–2434, 2014.

- [110] Yang Wu, Jun Wang, Chun Chen, Chan-Juan Liu, Feng-Ming Jin, and Ni Chen. Adaptive weighted gerchberg-saxton algorithm for generation of phase-only hologram with artifacts suppression. *Optics express*, 29(2):1412–1427, 2021.
- [111] M. O. Brown, T. Thiele, C. Kiehl, T.-W. Hsu, and C. A. Regal. Gray-molasses optical-tweezer loading: Controlling collisions for scaling atom-array assembly. *Physical Review X*, 9:011057, Mar 2019.
- [112] Adam L. Shaw, Pascal Scholl, Ran Finklestein, Ivaylo S. Madjarov, Brandon Grinkemeyer, and Manuel Endres. Dark-state enhanced loading of an optical tweezer array. *Physical Review Letters*, 130:193402, May 2023.
- [113] Daniel Barredo, Sylvain de Léséleuc, Vincent Lienhard, Thierry Lahaye, and Antoine Browaeys. An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science*, 354(6315):1021–1023, 2016.
- [114] Donggyu Kim, Alexander Keesling, Ahmed Omran, Harry Levine, Hannes Bernien, Markus Greiner, Mikhail D. Lukin, and Dirk R. Englund. Large-scale uniform optical focus array generation with a phase spatial light modulator. *Optics Letters*, 44(12):3178–3181, Jun 2019.
- [115] Philip Thomas Starkey. *A software framework for control and automation of precisely timed experiments*. PhD thesis, Monash University, Australia, 2019.
- [116] Andreas Trenkwalder, Matteo Zaccanti, and Nicola Poli. A flexible system-on-a-chip control hardware for atomic, molecular, and optical physics experiments. *Review of Scientific Instruments*, 92(10), 2021.
- [117] Omar Morandi, Sara Nicoletti, Vladislav Gavryusev, and Leonardo Fallani. Optimal control in phase space applied to minimal-time transfer of thermal atoms in optical traps. *Physical Review A*, 111(6):063312, 2025.
- [118] Andrea Fantini. Study and optimization of individual atom manipulation in quantum computational arrays, 2022.

- [119] D.M. Stamper-Kurn W. Ketterle, D.S. Durfee. Making, probing and understanding bose-einstein condensates. In S. Stringari M. Inguscio and C.E. Wieman, editors, *Bose-Einstein condensation in atomic gases*, volume 140 of *Proceedings of the International School of Physics "Enrico Fermi", Course CXL*, pages 67–176. IOS Press, Amsterdam, 1999.
- [120] Jonatan Hörschele, Sandra Buob, Antonio Rubio-Abadal, Vasiliy Makhalov, and Leticia Tarruell. Atom-number enhancement by shielding atoms from losses in strontium magneto-optical traps. *Physical Review Applied*, 19(6):064011, Jun 2023.
- [121] Martin Miranda, Ryotaro Inoue, Yuki Okuyama, Akimasa Nakamoto, and Mikio Kozuma. Site-resolved imaging of ytterbium atoms in a two-dimensional optical lattice. *Physical Review A*, 91(6):063414, 2015.
- [122] Karl D. Nelson, Xiao Li, and David S. Weiss. Imaging single atoms in a three-dimensional array. *Nature Physics*, 3(8):556–560, 2007.
- [123] Maxwell F. Parsons, Florian Huber, Anton Mazurenko, Christie S. Chiu, Widagdo Setiawan, Katherine Wooley-Brown, Sebastian Blatt, and Markus Greiner. Site-resolved imaging of fermionic ^6Li in an optical lattice. *Physical Review Letters*, 114:213002, May 2015.
- [124] G. J. A. Edge, R. Anderson, D. Jervis, D. C. McKay, R. Day, S. Trotzky, and J. H. Thywissen. Imaging and addressing of individual fermionic atoms in an optical lattice. *Physical Review A*, 92:063406, Dec 2015.
- [125] Alexandre Cooper, Jacob P. Covey, Ivaylo S. Madjarov, Sergey G. Porsev, Marianna S. Safronova, and Manuel Endres. Alkaline-earth atoms in optical tweezers. *Physical Review X*, 8:041055, Dec 2018.
- [126] Lukas Novotny and Bert Hecht. *Principles of nano-optics*. Cambridge university press, 2012.
- [127] Loïc Anderegg, Lawrence W. Cheuk, Yicheng Bao, Sean Burchesky, Wolfgang Ketterle, Kang-Kuen Ni, and John M. Doyle. An optical tweezer array of ultracold molecules. *Science*, 365(6458):1156–1158, 2019.

- [128] Martin Tom Lichtman. *Coherent operations, entanglement, and progress toward quantum search in a large 2D array of neutral atom qubits*. PhD thesis, The University of Wisconsin-Madison, 2015.
- [129] HF Arnoldus and G Nienhuis. Conditions for sub-poissonian photon statistics and squeezed states in resonance fluorescence. *Optica Acta: International Journal of Optics*, 30(11):1573–1585, 1983.
- [130] B. Efron and R.J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 1993.
- [131] Niamh Jackson, Ryan Hanley, Matthew Hill, Frédéric Leroux, Charles Adams, and Matthew Jones. Number-resolved imaging of $\{88\}$ sr atoms in a long working distance optical tweezer. *SciPost Physics*, 8(3):038, 2020.
- [132] Ralf Brooks et al. *Control and collisions of 87Rb and 133Cs atoms in optical tweezers*. PhD thesis, Durham University, 2022.
- [133] Lukas Homeier, Timothy J Harris, Tizian Blatz, Sebastian Geier, Simon Hollerith, Ulrich Schollwöck, Fabian Grusdt, and Annabelle Bohrdt. Antiferromagnetic bosonic t-j models and their quantum simulation in tweezer arrays. *Physical Review Letters*, 132(23):230401, 2024.
- [134] Mu Qiao, Gabriel Emperauger, Cheng Chen, Lukas Homeier, Simon Hollerith, Guillaume Bornet, Romain Martin, Bastien Gély, Lukas Klein, Daniel Barredo, et al. Realization of a doped quantum antiferromagnet in a rydberg tweezer array. *Nature*, 644(8078):889–895, 2025.
- [135] Max L Prichard, Benjamin M Spar, Ivan Morera, Eugene Demler, Zoe Z Yan, and Waseem S Bakr. Directly imaging spin polarons in a kinetically frustrated hubbard system. *Nature*, 629(8011):323–328, 2024.
- [136] Gabriel Emperauger, Mu Qiao, Cheng Chen, Filippo Caleca, Saverio Bocini, Marcus Bintz, Guillaume Bornet, Romain Martin, Bastien Gély, Lukas Klein, Daniel Barredo, Shubhayu Chatterjee, Norman Y. Yao, Fabio Mezzacapo, Thierry Lahaye, Tommaso Roscilde, and Antoine Browaeys. Tomonaga-luttinger liquid behavior in a rydberg-encoded spin chain. *Physical Review X*, 15:031021, Jul 2025.