

Rydberg Polariton Dynamics and Light Propagation in Ultracold Ytterbium Ensembles

Anthea Christina Nitsch

Masterarbeit in Physik
angefertigt im Institut für Angewandte Physik

vorgelegt der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität
Bonn

September 2026

I hereby declare that this thesis was formulated by myself and that no sources or tools other than those cited were used.

Bonn, 30.09.2026
.....
Date

.....
Signature

Referee and Supervisor: Prof. Sebastian Hofferberth
Referee: Prof. Daqing Wang

In the preparation of this thesis, digital tools—including AI-based language models—were used to support tasks such as code development for data analysis and visualization, debugging, $\text{\LaTeX}$ formatting assistance, literature research, and language refinement. All scientific content, numerical models, and results were independently designed, implemented, and verified by the author (see Table 1).

Table 1: Overview of digital and AI-based tools, structured by purpose of use.

Task / Purpose	Tools Used	Role of the Tools
Coding and Debugging	Claude (Anthropic)	Assistance with writing and debugging Python code for data plotting, fitting, and analysis; code documentation and structuring.
$\text{\LaTeX}$ Formatting and Typesetting	Claude (Anthropic)	Support with $\text{\LaTeX}$ syntax, table and figure formatting, and general document structuring.
Literature Research	Claude (Anthropic)	Assistance in locating and summarizing relevant literature on selected topics.
Writing and Language Polishing	Claude (Anthropic), LanguageTool (LTeX+)	Grammar and spelling checks, sentence clarity suggestions, and text restructuring.
Plotting of Data	Matplotlib (Python)	Generation and styling of data plots.
Data Analysis	SciPy (Python)	Fitting of models to experimental data.

Contents

1	Introduction	1
2	Apparatus for Dense Ytterbium Rydberg Ensembles	3
2.1	Atom Preparation and Trapping	3
2.1.1	Motorized Mirrors for Dipole Trap Alignment	6
2.2	Rydberg Excitation Scheme	9
2.3	Probe and Control Setup	10
2.3.1	Determining the Probe Beam Focus	10
2.3.2	Probe Beam Shape Scan Using Green Magneto-Optical Trap	11
2.3.3	Probe Beam Shape Scan Using Dipole Trap Position-Dependent Optical Depth	14
2.3.4	Calibrating the Probe Beam Acousto-Optic Modulator	17
2.3.5	Flat-Top Control Beam	19
2.4	Imaging of the Atomic Cloud and Detection of Photons and Ions	20
3	Light Propagation in Dense Atomic Media	23
3.1	Optical Depth in Atomic Media	24
3.2	Angular Spectrum Method	24
3.3	Validation of the Angular Spectrum Method	26
3.3.1	Comparison with the Thin-Lens Approximation	26
3.3.2	Recovery of the Lambert–Beer Law	28
3.4	Optical Depth for Different Atomic Cloud Configurations	30
3.4.1	Simulated and Measured Transmission Valley	33
3.4.2	Influence of Atomic Cloud Geometry	35
3.4.3	Influence of Atomic Conditions	38
4	Coherence and Lifetime of Collective Ytterbium Rydberg States	41
4.1	Rydberg State Lifetime	42
4.1.1	Measurement and Comparison to Theory Predictions	43
4.2	Coherence Times	46
4.2.1	Rydberg Collective State Coherence	46
4.2.2	Dephasing Mechanisms of the Rydberg Ensemble	50
4.2.2.1	Rydberg–Rydberg Interaction	50
4.2.2.2	Ion Creation	51
4.2.2.3	Longitudinal Dephasing and Transverse Mode Overlap	51
4.2.2.4	Collision-Induced Dephasing	55

4.2.3	Contribution of Rydberg Dimers	59
5	Two-Photon Correlations	62
5.1	Rydberg Blockade and Photon Antibunching	62
5.2	Second-Order Correlation Function	63
5.3	Experimentally Determined Second-Order Correlation Function	64
5.4	Ion Simulation	67
6	Conclusion	71
A	Supplementary Material	74
A.1	Probe Beam Tracing	74
A.2	Black Body Radiation and Spontaneous Emission Rates	75
A.3	Collisional Dephasing Maps for $n = 42, 73$, and 96	77
A.4	Ion Simulation	77
	Bibliography	81

Introduction

Single photons are robust carriers of quantum information, as they interact only weakly with their environment. Furthermore, they can be generated, manipulated, and detected with high efficiency [1]. This makes them a natural building block for quantum communication [2] and photonic quantum computation [1], where photons can also be stored in quantum memories [3]. However, the same weak coupling also suppresses interactions between photons. Single photons in conventional media as well as in free space interact only negligibly. In a medium, an interaction between photons requires an optical nonlinearity, i.e., a response of the medium that depends on the number of photons present, but the nonlinearity of conventional materials is negligible at the level of single photons [4]. Hence, facilitating strong photon–photon interactions, as required, e.g., for photonic quantum gates [5, 6], is a long-standing challenge in the field of quantum optics [4]. Strong single-photon nonlinearities have been realized, e.g., with single atoms in optical cavities [7] and with quantum dots coupled to nanophotonic waveguides [8]. In both cases, a single emitter can absorb only one photon at a time, so that the transmission of a second photon depends on the presence of the first.

A further approach, which is followed in this thesis, is to make photons interact effectively via the strong interactions between Rydberg atoms [9]. Such Rydberg-mediated interactions, which can be dissipative or dispersive depending on the scheme [9], have been applied, for example, to photonic phase gates [6], single-photon sources [10, 11], and single-photon transistors [12] and absorbers [13]. In this thesis, a resonant scheme based on [Electromagnetically Induced Transparency \(EIT\)](#) is used. A weak probe field, whose photons are to interact, couples an atomic ground state to an intermediate state, which is coupled to a Rydberg state by a strong control field. Under these conditions, the probe photons propagate through the atomic medium as slow, hybrid light–matter excitations known as Rydberg polaritons [14]. Due to the strong, long-range van der Waals interaction between Rydberg atoms [15], a Rydberg excitation prevents further excitations within the so-called blockade radius, so that the [EIT](#) condition is broken there for any further photon. A single photon thereby changes the transmission of the medium for subsequent photons within the blockade radius, which results in an effective, dissipative interaction between photons.

The properties of such a Rydberg medium are determined by the atomic species used. Most experiments in Rydberg quantum optics to date have been performed with alkali atoms, in particular [Rubidium-87 \(Rb-87\)](#) [6, 12, 16, 17]. In this thesis, [Ytterbium-174 \(Yb-174\)](#) is used instead, which offers three main advantages over alkali atoms: First, its second valence electron can be addressed optically while the first one occupies a Rydberg state, which provides additional means for the optical control and detection of

Rydberg excitations [18, 19]. Second, it possesses the narrow $^1S_0 \rightarrow ^3P_1$ intercombination line, which enables laser cooling to low temperatures in a two-color **Magneto-Optical Trap (MOT)** [20]. Third, the probe and control transitions at 399 nm and 395 nm are close in wavelength, which reduces the motional dephasing caused by the thermal motion of the atoms in a counter-propagating configuration. The experiments are performed in the Rydberg apparatus of the **Nonlinear Quantum Optics (NQO)** group at the University of Bonn [20], which, in addition to the probe and control lasers, comprises in-vacuum electrodes and a **Microchannel Plate (MCP)** detector for ion steering and detection, as well as a setup for detecting the photons transmitted through the atomic cloud.

In this thesis, three aspects are addressed that are central for the use of **Yb-174** atoms in this apparatus as a nonlinear medium for photons:

1. the determination of the **Optical Depth (OD)** of the dense atomic cloud, which sets the strength of the Rydberg-mediated nonlinearity [21],
2. the Rydberg state lifetime and the coherence time of collective Rydberg excitations, which determine how long photons can be stored, and the mechanisms limiting it in the dense **Yb-174** ensemble, and
3. the effective photon–photon interaction mediated by the Rydberg blockade, which manifests itself as antibunching of the transmitted probe light.

In chapter 2, the apparatus is described. The contributions to the apparatus made in the course of this thesis are the automation of the dipole trap alignment using motorized mirrors, two independent methods to precisely locate the probe beam focus, and the calibration of the probe pulse shape.

In chapter 3, a beam propagation simulation based on the **Angular Spectrum Method (ASM)** is introduced and validated against known analytical results. The density of the cloud varies transversally, so that the probe beam is diffracted and refracted. As a result, the conventional one-dimensional model is shown to systematically overestimate the **OD** experienced by the probe beam.

In chapter 4, it is investigated how long a Rydberg polariton survives as a coherent collective excitation in the atomic ensemble. The Rydberg state lifetime is measured and compared to theoretical predictions [22, 23], which bounds the coherence time from above. The coherence time of the stored collective excitation is then compared to this bound. Of the several dephasing mechanisms examined, collisions between Rydberg and ground-state atoms are identified as the dominant one, shortening the coherence time well below the Rydberg lifetime. Further, a periodic coherence revival attributed to the excitation of Rydberg dimers is observed.

Finally, in chapter 5, the photon statistics of the transmitted probe light are examined. The two-photon correlation function is presented. Alongside the expected antibunching dip, an additional structure is observed. It is attributed to the ionization of ground-state atoms by the control light and is examined with a minimal simulation.

Apparatus for Dense Ytterbium Rydberg Ensembles

Enabling experiments in Rydberg quantum optics places several requirements on an apparatus. Since the strength of the Rydberg-mediated nonlinearity scales with the [OD](#) per blockade radius, a cold and dense atomic ensemble has to be prepared and overlapped precisely with a tightly focused probe beam. The probe pulses have to be shaped precisely and switched off completely, while the strong control field must not push the atoms out of the trap. Since Rydberg states are highly sensitive to electric fields, stray fields at the position of the atoms have to be compensated. Finally, both the transmitted photons and the ions created from Rydberg atoms have to be detected with high temporal resolution. In the [Ytterbium Quantum Optics \(YQO\)](#) experiment, these requirements are met for [Yb-174](#).

In section [2.1](#), the preparation of the atoms in a two-color [MOT](#) and their transfer into an optical dipole trap are described, together with motorized mirrors implemented for the alignment of the dipole trap. The two-photon Rydberg excitation scheme and the Rydberg blockade are introduced in section [2.2](#). The probe and control setup is presented in section [2.3](#). Two methods are developed there to determine the probe beam focus, and the probe pulse shape is calibrated. In addition, a flat-top control beam is implemented, which reduces atom loss. Finally, the photon and ion detection setups are discussed in section [2.4](#).

2.1 Atom Preparation and Trapping

The characterization measurements in chapters [3](#), [4](#), and [5](#) require a cold and dense atomic ensemble that is positioned precisely at the probe beam focus. In this section, the preparation of such an ensemble and the automation of the alignment of the dipole trap beams are described.

For this purpose, [Yb-174](#) atoms are first loaded from dispensers into a two-dimensional [MOT](#) operating on the broad $^1S_0 \rightarrow ^1P_1$ transition with wavelength 399 nm. A weak, resonant laser beam on the same transition, the so-called push beam, accelerates slowly moving atoms along the axis of the two-dimensional [MOT](#) to the minimum velocity required to traverse the differential pumping stage. This increases the number of atoms delivered to the main vacuum chamber, in which all experiments are performed. After propagating through the differential pumping stage, the atoms are captured by a two-color three-dimensional [MOT](#), initially operating on the blue $^1S_0 \rightarrow ^1P_1$ transition. During this

blue-MOT stage, a fraction of the atoms is optically shelved into the long-lived $^3P_1$ triplet state using positively detuned green light. This suppresses atom loss from light-assisted collisions on the strong blue transition and thus increases the number of retained atoms. Subsequently, the atoms are further cooled using the $^1S_0 \rightarrow ^3P_1$ intercombination line, reaching temperatures below $10\text{ }\mu\text{K}$ with typical atom numbers of around $7 \cdot 10^6$. A more detailed description of the atom cooling scheme using the combination of two-dimensional and two-color three-dimensional MOTs can be found in [20, 24].

Following the green MOT stage, the atom cloud is spatially compressed to a smaller size and positioned so as to maximize the transfer efficiency into the optical trap. The compression and repositioning are done by increasing the magnetic field gradient and shifting the cloud with magnetic offset fields. After cooling the atoms with the MOT configuration described above, they are loaded into an optical dipole trap formed by two crossed focused Gaussian beams at a wavelength of 1064 nm . The two beams cross at a half-angle of 10.12° , with a power of up to 16.0 W and a beam waist of $42.6\text{ }\mu\text{m}$ in each beam. The atom number and the temperature of the atoms trapped in the dipole trap can be adjusted in two ways:

- the number of atoms loaded into the dipole trap is controlled via the loading time of the blue three-dimensional MOT, and
- the density and temperature of the atoms are varied via the dipole trap power.

To reach low temperatures, the dipole trap beam power can be adiabatically decreased to a lower intensity, to remove the hotter atoms from the dipole trap potential, and then increased again, to recompress the remaining atoms.

The resulting cold and dense atomic cloud forms the atomic medium, to which the excitation laser pulses are applied. The complete experimental cycle is shown in figure 2.1. First, N signal pulses are applied to the cloud, in which the probe and/or the control laser are switched on with pulse shapes that depend on the measurement. During these pulses and the photon detection and ion extraction (see section 2.4), the dipole trap is switched off. The signal pulses are followed by N reference pulses with identical excitation laser pulses, during which the dipole trap is switched off, so that no atomic cloud is present.




Figure 2.1: Experimental cycle. The bottom row shows the full sequence: preparation of the atomic cloud in the blue and green MOT, loading and holding in the dipole trap, followed by N signal pulses and, after a delay, N reference pulses without atoms. The top panel shows the logic levels during one repetition: the dipole trap and the Rydberg excitation lasers, the gate of the Single-Photon Counting Modules (SPCMs), and the ion detection system consisting of high voltages for ionization and steering and a MCP detector. Adapted from [25].

2.1.1 Motorized Mirrors for Dipole Trap Alignment

In the crossed dipole trap, the atoms are trapped at the intersection of the two beams. After the initial alignment, this position must be periodically readjusted, since thermal and mechanical drifts cause a shift over time. This was originally done by manually tuning two mirrors.




Figure 2.2: Probe and control setup as well as dipole trap beam configuration. The atomic cloud is trapped in a crossed dipole trap at the probe and control beam foci.

As shown in figure 2.2, the intersection of the two dipole trap beams, and thus the atomic cloud, has to coincide with the focus of the probe beam. The two dipole trap beams are orthogonally polarized to prevent interference between them. The probe and control beams propagate through the main vacuum chamber in opposite directions with the same circular polarization and are separated by dichroic filters before and after the chamber. These filters prevent the strong control light and back-reflected probe light from entering the collection fiber, which guides the transmitted probe light to the photon detection setup.

Aligning the crossed dipole trap beams to the probe beam focus by hand was labor-intensive. Moreover, the achievable alignment accuracy was limited to about $15\text{ }\mu\text{m}$. This exceeds the waists of the probe and control beams of $5\text{ }\mu\text{m}$ and $10\text{ }\mu\text{m}$, respectively. With an alignment precision of this order, it is difficult to position the trap such that the probe beam passes through the densest part of the cloud, which reduces the OD experienced by the probe light. In addition, the repeated manual access to the beam path posed an increased laser safety risk. To address these issues, the two mirrors marked in figure 2.2 were equipped with motorized actuators. These motorized mirrors consist of a high-stability mirror mount, to which two linear actuators are attached, one for each tilt axis of the mirror.¹ In the experiment, these actuators are addressed via the Physik Instrumente (PI) drivers², which are in turn controlled by software provided by PI. A detailed characterization of identical actuators, including their working principle, can be found in [26].

To characterize the achievable step size and its reproducibility for these linear actuators in the chosen mirror mounts, a test setup was built, shown in figure 2.3. A laser beam is reflected by a motorized

¹ Mount: Radiant Dyes high-stability mount. Actuators: PiezoMike Linear Actuators, Model N-470, rebranded by Radiant Dyes as N-970.

² E-870.4G PiShift drivers

mirror and focused by a lens onto a camera³ placed in the focal plane of the lens. The beam position on the camera is extracted from each recorded image by independently fitting a Gaussian function along both camera axes to the beam spot; the peak position of this Gaussian function is taken as the beam position, with its uncertainty given directly by the standard error of the extracted peak position.




Figure 2.3: Test setup for characterization of the motorized mirrors. A laser beam coming from a fiber is reflected by the motorized mirror being tested, focused by a lens, and recorded by a camera in the focal plane with fine enough resolution to resolve the beam focus. With this setup, it is possible to alternately tilt the mirror vertically and horizontally for a set number of steps with varying numbers of increments per step. The two graphs on the right show the recorded beam positions for 1 increment (left) and 100 increments (right) per step, with the color of each data point denoting its measurement number.

Here, an increment denotes the smallest motion of the actuator, and a step consists of a set number of increments. The mirror is alternately tilted by a fixed number of steps horizontally and then by the same number of steps vertically, and this sequence is repeated once in the opposite directions, so that the beam ideally traces out a closed square loop and returns to its starting position. The resulting beam positions are shown in figure 2.3 for 1 and 100 increments per step, with the driver parameters listed in table 2.1.

Table 2.1: PIShift driver parameters used to operate the linear actuators during characterization.

Parameter	Value	Description
PIShift Upper Supply Voltage	80.0 V	Upper voltage limit of the piezo drive signal
PIShift Lower Supply Voltage	0.0 V	Lower voltage limit of the piezo drive signal
PIShift Forward Current	0.14 A	Charging current during forward motion
PIShift Backward Current	−0.11 A	Charging current during backward motion
PIShift Frequency	400.0 Hz	Operating frequency of the step-mode drive signal
PIShift Charge Cycle	1.0	Duty cycle of the current source per drive signal period

For both step sizes, the beam does not return exactly to its starting point. Thus, for a fixed number of steps, the mirror covers a slightly different angle in the forward than in the backward direction along the same axis. This asymmetry is not unexpected, since the charging currents during forward and backward motion already differ to compensate for mechanical or friction asymmetries in the stick-slip mechanism of the actuators (see table 2.1). For single increments, the data points additionally scatter perpendicular to the direction currently being stepped, by approximately 0.3 to 0.7 μm . This scatter exceeds the typical

³ MT9J001 Arducam

statistical uncertainty of $0.08 \mu\text{m}$ on the extracted peak position and thus indicates the precision that the mirror can achieve.

The beam position on the camera depends on the focal length of the lens used and thus on the specific test setup. A quantity independent of the lens is the tilt angle of the mirror. Since the camera is placed in the focal plane of the lens, all rays with the same angle are focused onto the same point, such that the beam position on the camera depends only on the deflection angle of the beam and not on the distance between mirror and lens. As the beam is deflected by twice the mirror tilt, the angle θ can be determined to be [26]

$$\theta = \frac{\arctan(\Delta x/f)}{2} \approx \frac{\Delta x}{2f}, \quad (2.1)$$

where Δx is the deflection of the beam from the camera center and $f = 250 \text{ mm}$ is the focal length of the lens focusing the beam onto the camera. The resulting mean beam displacement and tilt angle per increment are listed in table 2.2 for both mirrors, both tilt axes, and both directions. The increment size of both mirrors lies between about $0.5 \mu\text{rad}$ and $0.9 \mu\text{rad}$ per increment, corresponding to 0.24 to $0.45 \mu\text{m}$ per increment on the camera. This is consistent with the nominal step size of $0.67 \mu\text{rad}$ and the values of 0.53 to $1.00 \mu\text{rad}$ measured for identical actuators in [26]. The forward–backward asymmetry observed above is pronounced to varying degrees depending on the mirror and the tilt axis.

Table 2.2: Mean angular and positional beam deflection per increment, averaged over all tested numbers of increments per step (1, 10, 100, and 200), for forward and backward motion in the horizontal (x) and vertical (y) tilt directions, shown separately for both mirrors.

Mirror	Forward	Backward
Mirror 1, Horizontal Angle ($\mu\text{rad}/\text{increment}$)	0.7207 ± 0.0005	0.8982 ± 0.0012
Mirror 1, Vertical Angle ($\mu\text{rad}/\text{increment}$)	0.4793 ± 0.0004	0.4910 ± 0.0003
Mirror 1, Horizontal Position ($\mu\text{m}/\text{increment}$)	0.3632 ± 0.0003	0.4507 ± 0.0006
Mirror 1, Vertical Position ($\mu\text{m}/\text{increment}$)	0.2395 ± 0.0002	0.2455 ± 0.0002
Mirror 2, Horizontal Angle ($\mu\text{rad}/\text{increment}$)	0.7280 ± 0.0018	0.6359 ± 0.0034
Mirror 2, Vertical Angle ($\mu\text{rad}/\text{increment}$)	0.5421 ± 0.0003	0.5415 ± 0.0006
Mirror 2, Horizontal Position ($\mu\text{m}/\text{increment}$)	0.3727 ± 0.0009	0.3187 ± 0.0014
Mirror 2, Vertical Position ($\mu\text{m}/\text{increment}$)	0.2715 ± 0.0001	0.2708 ± 0.0003

In practice, the asymmetry of the increment size is accounted for by cross-checking the position of the atomic cloud in the dipole trap with absorption imaging (see section 2.4). The positioning accuracy is obtained by comparing the measured cloud position before and after adjusting the mirror tilt by a single increment, or, for the manual case, by hand as gently as possible. With this procedure, the positioning accuracy is improved from about $15 \mu\text{m}$ with manual alignment to about $2 \mu\text{m}$ with the motorized mirrors, which lies well below the beam waists of 5 to $10 \mu\text{m}$. In addition, manual access to the beam path is no longer required, and the time and effort for each realignment are substantially reduced.

2.2 Rydberg Excitation Scheme

After the atoms are trapped in the dipole trap, they are excited to either the intermediate or the Rydberg state. A direct single-photon transition from the ground state $|g\rangle$ to the Rydberg state $|r\rangle$ is dipole-forbidden, since both are S states. Instead, a two-photon excitation scheme via the intermediate state $|e\rangle$ is used, which not only makes the transition dipole-allowed and relaxes the power requirement, but also provides the three-level structure needed to access further interesting Rydberg physics, such as EIT.




Figure 2.4: Rydberg excitation and MOT cooling scheme for Yb-174. The Rydberg state $|r\rangle$ is excited from the ground state $|g\rangle$ via an intermediate state $|e\rangle$, with the excitation lasers detuned by the probe detuning Δ_p and the two-photon detuning δ . The blue MOT (2D and 3D) uses the same transition as the probe beam, while the green MOT uses the $6s^2 1S_0 \rightarrow 6s6p^3 P_1$ transition.

The excitation scheme is shown in figure 2.4. To excite atoms to the Rydberg state, a weak probe laser with wavelength $\lambda_p \approx 399$ nm drives the $6s^2 1S_0 \rightarrow 6s6p^1 P_1$ transition, detuned by Δ_p from the intermediate state $|e\rangle$, and a strong control laser, $\lambda_c \approx 395$ nm, drives the $6s6p^1 P_1 \rightarrow 6sns^1 S_0$ transition, detuned by Δ_c from the Rydberg state $|r\rangle$. Together, these give the two-photon detuning $\delta = \Delta_p + \Delta_c$. Because the transition matrix element to the Rydberg state scales unfavorably with the effective principal quantum number ν , the control transition requires comparatively high power. The probe, by contrast, is deliberately kept weak to remain in the few-photon regime required for the photon-photon interaction experiments this apparatus targets.

Once excited to the Rydberg state, atoms interact via long-range van der Waals forces, owing to their large electric dipole moment, $V_{\text{vdW}}(r) = -C_6/r^6$, with C_6 scaling as $C_6 \propto \nu^{11}$ with the effective principal quantum number ν [15]. Within a certain interaction range – the Rydberg blockade radius r_b – this interaction shifts a second atom’s Rydberg level out of resonance with the excitation lasers, an effect known as the Rydberg blockade. The blockade radius follows from the condition that this interaction shift equals twice the linewidth of the excitation process, $V_{\text{vdW}}(r_b) = 2\hbar\Gamma_{\text{EIT}}$, giving

$$r_b = \left(\frac{C_6}{2\hbar\Gamma_{\text{EIT}}} \right)^{1/6}, \quad (2.2)$$

where the [EIT](#) linewidth $\Gamma_{\text{EIT}} = \Omega_c^2/4\gamma_{ge}$ is set by the control Rabi frequency Ω_c and the decoherence rate γ_{ge} of the probe transition, which is given by $\gamma_{ge} = \Gamma/2$ for a purely radiatively broadened transition with the natural linewidth Γ . Within this range, only a single Rydberg excitation can be created, and thus only a single photon can propagate as a Rydberg polariton, while a second photon is scattered.

2.3 Probe and Control Setup

The probe beam focus serves as the fixed reference position of the experiment: the dipole trap is aligned to it, and the control beam is then aligned to the resulting position of the atomic cloud (see figure [2.2](#)). In this section, two methods to determine the probe beam focus are presented, followed by the calibration of the probe pulse shape and the implementation of a flat-top control beam.

2.3.1 Determining the Probe Beam Focus

The dipole trap must be finely aligned to a chosen reference position, which in this case is the probe beam focus. Aligning the atomic cloud to the probe beam focus offers two advantages. First, a sufficiently small beam waist allows the entire transverse extent of the beam to lie within a single blockade radius for high enough Rydberg states. As a result, only one Rydberg excitation – and thus only one photon propagating as a Rydberg polariton – can be created at a time across the beam. Along the cloud, which is longer than the blockade radius, several excitations can still be present, spaced by at least the blockade radius. This is essential for the single-photon and photon–photon interaction measurements this apparatus targets. Second, the coupling between the probe field and the atoms is maximized at the focus (see section [2.2](#)).

After aligning the atoms to the probe beam focus, the control beam is in turn aligned to the atomic cloud. To perform this fine alignment, the probe beam focus position has to be precisely determined. In the directions perpendicular to the probe beam propagation direction, this can be achieved by splitting the atom cloud captured by the dipole trap with a high-power probe beam, which locally displaces atoms via the scattering force it exerts on them (see section [2.3.2](#)), creating a visible hole in the cloud whose transverse position marks where the probe beam intersected it. This cloud shape is recorded using the same absorption imaging technique described in section [2.4](#): an imaging beam is passed through the cloud and the transmitted light recorded on a camera, from which the atom distribution is extracted, directly revealing the probe beam’s transverse center.

In the axial (probe beam propagation) direction, determining the focus position is not as straightforward, and different methods were explored. In the course of this thesis, two such methods were developed and implemented: the first exploits the scattering force the probe beam exerts on atoms held in the green [MOT](#) to trace out the beam’s profile at different points along its propagation axis; the second instead uses the dipole-trap-position-dependent [OD](#) of atoms illuminated by the probe to the same end. Both methods, and the resulting focus positions, are described in detail in sections [2.3.2](#) and [2.3.3](#), respectively.

2.3.2 Probe Beam Shape Scan Using Green Magneto-Optical Trap

To determine the probe beam's spatial profile, the scattering force it exerts on the atoms released from the green MOT is used. Since the probe beam is resonant with an atomic transition, each photon absorption and re-emission cycle transfers the photon recoil momentum to the atoms. Over many cycles, this accumulates into a net scattering force that displaces the atoms along the probe beam, so that they are shifted with respect to the rest of the cloud. Strictly speaking, the measurement below maps this atomic response rather than the probe beam directly.

Where the probe intensity saturates the atomic transition, the scattering force and thus the atomic displacement are likewise saturated and not proportional to the local probe intensity, instead approaching a constant, intensity-independent value [27, 28]. The atoms displaced from the probe beam path leave a line of reduced OD in the absorption image (see figure 2.6). Since the region of saturated scattering force is narrowest where the beam is narrowest, the probe beam focus can still be determined, while the extracted width is larger than the actual beam radius.

Implementing this concept experimentally required delicate tuning of the delay between turning on the probe beam and imaging the cloud, as explained below. Further, the cloud shape in the green MOT is not perfectly Gaussian and is slightly position dependent, as observed in the reference absorption images taken without the probe beam, making it difficult to extract the probe beam shape cleanly. In the following, the sequence that gives a good impression of the probe beam shape is briefly explained; an example sequence is shown in figure 2.5.




Figure 2.5: Timing of the probe beam, imaging, and green MOT to achieve the probe beam trace measurement. The green MOT is turned off, and, with a delay, the probe beam and imaging are turned on, with the imaging happening at the same time as the probe beam or slightly delayed (here at the same time).

First, the MOT is turned off and the atoms are left to expand for a variable time until the shape of the atom cloud is approximately Gaussian: during this free expansion, the spatial distribution of the cloud is increasingly determined by the Gaussian thermal velocity distribution of the atoms, so that irregularities of the initial MOT shape are washed out. In the shown sequence, the MOT expansion time is roughly 2 ms. Following this, the probe beam, with a power of around 2 mW, is turned on for 40 μ s and creates a hole in the cloud by pushing the atoms along its path via the scattering force (see figure 2.6). Since the atoms are no longer confined by the MOT at this point, it is crucial to find a combination of probe beam

power and small-to-no imaging delay between turning on the probe beam and imaging the cloud: the probe power has to be high enough to produce a signal that is clearly resolvable against the imaging noise, while the imaging delay has to be short enough that the atoms have not moved appreciably in the meantime beyond the displacement induced by the probe beam – such motion can shift the extracted probe beam focus position by several hundred micrometers.

The imaging delay in this sequence was fine-tuned by cross-checking against the absorption imaging; it was found that the effect of the probe beam on the atomic cloud was already visible without any additional delay in the sequence. Note that the times set in the sequence can differ from the actual switching times of the experiment's components, due to differing signal delays. Imaging the MOT with the hole created by the probe beam is repeated for varying positions of the green MOT along the probe beam path to reconstruct the probe beam's spatial profile along its full propagation direction. This is achieved by varying the magnetic offset fields that determine the position of the MOT. Note that there is no guarantee that the probe beam path and the spatial axes set by the magnetic field coils of the MOT align, meaning that more than one magnetic field must be adjusted to trace the probe beam: the MOT has to be moved along a combination of the coil axes, such that it follows the probe beam path.




Figure 2.6: Absorption images of the green MOT after an expansion time of about 2 ms, recorded along the vertical imaging axis, each averaged over 23 measurement cycles. (a) During the probe pulse, the probe beam, which propagates along x , appears as a narrow line of reduced OD through the cloud, since the atoms in its path are pushed along the probe beam. As the probe beam is much narrower than the cloud, only a small fraction of the atoms along the line of sight is affected, so that the line is shallow. (b) The reference image is taken without the probe beam and is subtracted from (a) to isolate the hole created by the probe beam.

To extract the probe beam shape, absorption images (see section 2.4) are recorded along both the vertical and horizontal imaging axes. For each imaging axis, a separate reference measurement of the MOT shape without the probe beam is taken at each position along the probe beam. The reference image is then subtracted from the image taken with the probe pulse, isolating the atomic displacement induced by the probe beam from the underlying MOT cloud shape, and the resulting traces from all positions are stacked on top of each other to form a single two-dimensional image, with position along the probe beam on one axis and the transverse position on the other.

It has to be noted that the probe beam trace measured here does not directly reflect the properties of the

probe beam, such as its waist. At the probe power used, the transition is saturated across essentially the whole beam profile, so that the scattering force and thus the atomic response are no longer proportional to the local probe intensity. The trace therefore maps the region in which the probe intensity exceeds the saturation intensity, which is wider than the beam itself. This nonlinearity can be seen in the upper plot in figure 2.7.




Figure 2.7: Vertical probe beam trace and slices transversal to the probe beam, showing the beam profile at varying points along its propagation direction. Using the method explained in this section, the probe beam trace is extracted for both vertical and horizontal imaging (see appendix A). The white dashed lines in the upper plot mark the positions of the transverse slices. In the lower plots, the averaged signal (blue circles) is shown together with the double-Gaussian model of equation 2.3 (black line), from which the beam width at each position is extracted.

From this plot, the probe beam shape can be extracted by taking traces perpendicular to the probe beam, indicated by three white dashed lines in the image. The lower plots show the profile of the probe beam along these three example perpendicular traces. The imaged probe beam is not necessarily very Gaussian in shape, which makes extracting the probe beam focus challenging. In particular, an additional peak appears next to the main peak, which is attributed to the reflection of the probe beam at the front and back surface of the dichroic mirror (see figure 2.2), producing a second, weaker, and laterally displaced copy of the probe beam. To best model the data, a sum of two Gaussians with an offset is used instead,

$$f(y) = a_1 \exp\left(-\frac{(y - \mu_1)^2}{2\sigma_1^2}\right) + a_2 \exp\left(-\frac{(y - \mu_2)^2}{2\sigma_2^2}\right) + b. \quad (2.3)$$

With this, the double-Gaussian structure visible in the probe beam trace can be resolved, and the true peak of the probe beam trace can be extracted. Repeating this for various positions along the probe beam, the extracted beam width at each position is described by the standard Gaussian beam profile

$$w(x) = w_0 \sqrt{1 + \left(\frac{x - x_0}{z_R}\right)^2}, \quad (2.4)$$

where w_0 is the extracted waist, z_R the Rayleigh range, and x_0 the probe beam focus position, yielding a beam waist plot. This plot is shown in figure 2.8.




Figure 2.8: Extracted waist from the probe beam trace. The beam widths extracted at each position are shown as circles with error bars, and the Gaussian beam profile of equation 2.4 adapted to them as a black line.

Here, only an example measurement is shown. While testing this method, significantly more measurements were taken, and it was generally observed that the extracted probe beam waist varies from measurement to measurement even under nominally identical experimental conditions. This variability is caused by the difficulty of fitting the probe beam perpendicular slices, owing to irregularities in the beam shape as well as some variability in the MOT shape between the probe measurement and the reference measurement. Because of these added uncertainties, the error on the probe beam focus position should be estimated as larger than the statistical uncertainty shown in figure 2.8, at about $200\text{ }\mu\text{m}$. This estimated uncertainty is, in fact, of the same order of magnitude as the probe beam’s Rayleigh range. The extracted probe beam waist is larger than the beam’s known waist, due to the effects described above.

Nonetheless, this method determines the probe beam focus position with a useful level of precision. The probe beam focus was found to be at about $(1\,382 \pm 200)\text{ }\mu\text{m}$, sufficient for the alignment steps that follow. The relatively large uncertainty on the found position is caused by the large variation of the extracted probe beam waist between measurements. With the tools currently available in the experimental setup, this is the highest precision achievable. Because of this, a second method to determine the probe beam focus was implemented, described in detail in the following section 2.3.3.

2.3.3 Probe Beam Shape Scan Using Dipole Trap Position-Dependent Optical Depth

As an alternative to the probe beam trace measurement described in section 2.3.2, the probe beam focus can also be determined using the atoms held in the dipole trap itself, rather than the green MOT. Since the dipole trap can be repositioned with high precision using the motorized mirrors described in section 2.1.1, the trap can be scanned through the probe beam’s area while measuring the resulting OD, i.e., the attenuation of the probe light by the atoms (see equation 2.6), at each position. Turning on only the probe beam to excite the atoms to the intermediate state effectively creates a two-level system (see figure 2.4). To determine the probe beam’s spatial profile using the dipole trap, the attenuation the probe beam experiences when passing through the trapped atoms (see figure 2.2) is measured for varying positions of the dipole trap relative to the probe beam. This attenuation is quantified through the transmission T , defined as [9]

$$T = \frac{I_{\text{out}}}{I_{\text{in}}} = e^{-\text{OD}(\Delta_p)}, \quad (2.5)$$

where I_{in} and I_{out} are the incoming and outgoing beam intensities, respectively, and $\text{OD}(\Delta_p)$ is the probe-detuning-dependent OD. The OD at a given trap position is therefore extracted from the measured transmission as a function of probe detuning, as described in detail in section 3.1. The OD in an atomic medium is defined as

$$\text{OD} = \int \sigma_a(\Delta_p) \rho(x, y, z) dx, \quad (2.6)$$

where $\sigma_a(\Delta_p)$ is the atomic absorption cross section and $\rho(x, y, z)$ is the atomic density, with x taken along the probe beam's propagation direction. It has to be noted that the OD scales directly with the atom number as well as with the cloud shape. This comes from the atomic density, which can be rewritten as $\rho(x, y, z) = N_a \cdot f(x, y, z)$, where N_a is the total atom number and $f(x, y, z)$ is the normalized spatial distribution of the cloud. In this measurement, it is assumed that the cloud shape and size remain unchanged, while the atom number can change due to the variable loading efficiency of the green MOT into the dipole trap for varying dipole trap positions. This loading efficiency was optimized for one position and can decrease for other positions. The position of the dipole trap for each measurement is determined using a separate absorption image. As will be described in more detail in chapter 3, the OD determined with this method does not accurately reflect the OD experienced by the probe beam, due to diffraction and refraction of the probe light by the atomic cloud. For the purpose of comparing this quantity across varying trap positions, however, this is sufficient, and it will simply be referred to as the OD in this section.

Having determined the OD from the probe-detuning-dependent transmission of the atoms, a correlation between the position of the dipole trap relative to the probe beam path and the OD can be determined. The OD varies with this relative position. Using this dependence, the probe beam focus can be determined by measuring the transmission valley for varying dipole trap positions and determining the OD. Without the upgrade of the motorized mirrors described in section 2.1.1, this kind of measurement would not have been possible due to both the insufficient positioning precision and the time it would have taken to reposition the dipole trap manually at each new position. It is important to mention that the alignment of the two dipole trap beams relative to each other is not re-optimized at each new position, but it is assumed that the alignment is not significantly worse. Further, the atom number at each dipole trap position was measured to be able to mostly eliminate the dependence on the overall atomic conditions in the dipole trap. This follows directly from equation 2.6, since dividing the OD by N_a removes its explicit dependence on atom number. Dividing the OD by the atom number in the dipole trap at each position therefore yields a quantity that, for low probe power and small OD, tracks the local probe beam intensity convolved with the cloud shape, from which the probe beam shape can be inferred. This measurement can be seen in figure 2.9.

To extract the probe beam focus position, a model function is applied to the data plotted in this figure. It has to be kept in mind that what is measured here is not the probe beam shape directly, but a convolution of the Gaussian-shaped atomic cloud and the beam shape. Thus, a somewhat heuristic model was chosen to describe this shape in the form of




Figure 2.9: Vertical probe beam trace using the measured **OD** and atom number in the dipole trap. The measurement results are plotted as squares of varying sizes corresponding to their uncertainties. The color of the square denotes its value. The two-dimensional model function adapted to these data is shown in the background. The gray dashed lines mark the axes parallel and perpendicular to the probe beam’s propagation direction.

$$A \cdot \left(\frac{w_0}{w(x_{\text{probe}})} \right)^2 \cdot \exp \left(- \frac{2z_{\text{probe}}^2}{w(x_{\text{probe}})^2} \right), \quad (2.7)$$


where A is an amplitude, $w(x_{\text{probe}})$ is the beam width according to equation 2.4, and z_{probe} is the coordinate perpendicular to the probe beam. This model was chosen because it is the beam’s own intensity profile, without an explicit term accounting for the atomic cloud’s finite size. While convolving this profile with a Gaussian cloud, as discussed above, systematically biases the extracted waist w_0 and Rayleigh range z_R , it does not shift the position of the profile’s maximum along x . The extracted focus position x_0 can therefore be trusted even though w_0 and z_R cannot. Here, the extracted waist is defined according to equation 2.4 with the coordinates x_{probe} and z_{probe} . These coordinates are rotated by the angle θ with respect to the imaging axes and shifted such that x_{probe} runs parallel to the probe beam’s propagation direction and z_{probe} perpendicular to it, with the origin placed at the extracted probe beam focus. With this method, the probe beam focus is determined to be at $x_0 = (1\,278 \pm 11) \mu\text{m}$. The other extracted parameters, such as the probe beam waist and the Rayleigh range of the beam, should not be compared to the actual values of the probe beam, since again, what is measured here is not the direct probe beam shape but a convolution of probe beam shape and cloud shape. The extracted angle θ indicates that the probe beam axis is slightly rotated relative to the imaging axis.

Thus, this method yields a reasonable value for the probe beam focus position. Both methods using the green **MOT** and the dipole trap yield overlapping results with $(1\,382 \pm 200) \mu\text{m}$ and $(1\,278 \pm 11) \mu\text{m}$, respectively. The uncertainty on the position extracted from the **OD** in the dipole trap is the statistical uncertainty of the model parameter. This model probably underestimates the error due to the convolution of the probe beam and atom cloud shape. Because this measurement is highly time-consuming, it is impractical for fast determination of the probe beam focus; consequently, it was performed only once, and no estimate of the shot-to-shot fluctuation in the focus position is available. The two results nonetheless

agree within one Rayleigh range of the probe beam. Within this range, the beam width changes only slowly with position, making it difficult to determine the probe beam focus with any better precision than the Rayleigh range itself. The method based on the dipole trap yields an independent confirmation of the focus position, obtained directly with the trapped atoms.

2.3.4 Calibrating the Probe Beam Acousto-Optic Modulator

After determining the probe beam focus and aligning the dipole trap accordingly, precise control over the probe beam pulse shape becomes essential for the coherence time measurement, which determines the storage time of Rydberg polaritons (see chapter 4). The probe pulses are shaped by an **Acousto-Optic Modulator (AOM)** in the probe beam path, which diffracts part of the probe light into the fiber leading to the experiment, with a diffracted power that depends on the applied **Radio Frequency (RF)** signal. A calibration of the **AOM** is needed because the response of the **AOM** is nonlinear, which distorts any pulse shape applied as an **RF** signal. It is furthermore essential that only a negligible amount of residual probe light is emitted by the setup when the probe beam is switched off, since residual probe light can introduce an additional decoherence channel, artificially shortening the measured coherence times.



(a) **AOM** response curve. The **AOM** response is obtained by measuring the laser power after a fiber into which the -1 st diffraction order of the **AOM** beam was coupled, for a linearly rising **RF** signal. The black dashed lines show the region used to construct the calibration curve in figure 2.10(b). The **RF** signal is plotted in blue, the **AOM** response in purple.

(b) **AOM** calibration curve. In the time interval signified by the dashed lines in figure 2.10(a), the **RF** voltage is plotted against the normalized photodiode signal. The calibration curve (green) is extracted by taking a running mean over 30 data points and interpolating between the resulting points. The curve is clipped at 0 and 10 on the normalized photodiode signal axis.


Figure 2.10: **AOM** response and calibration.

This section introduces the probe beam **AOM** calibration used to generate the precise pulse shapes required for a successful coherence time measurement. To calibrate the **AOM**, its response curve first has to be recorded. This is done by sending a linearly ramped **RF** signal to the **AOM** and recording, on a photodiode, the diffracted output power coupled into a fiber after the **AOM**, across the full range of **RF** voltages. The measured calibration curve is then extracted from this response curve by rescaling both the photodiode signal and the **RF** voltage to the same normalized range (in this case the interval from 0 to 10, corresponding to the possible **RF** voltages in the experiment) and plotting the **RF** voltage against the

photodiode signal. This effectively inverts the **AOM** response, giving the **RF** voltage required for any desired diffracted power. Both curves are shown in figure 2.10.

In the range where the **AOM** response is monotonic, i.e., where each diffracted power corresponds to exactly one **RF** voltage – signified by the black dashed lines – an interpolating function is constructed from the measured calibration data to extract the **RF** voltage corresponding to every measured diffracted power. To obtain the best possible result, an interpolation of the running mean over 30 consecutive data points of the measured calibration curve is used. This interpolation was chosen because model functions such as a polynomial or a hyperbolic tangent could not follow the measured data at the edges of the calibration curve, where the slope becomes too steep. Furthermore, the calibration curve only spans values from around 1.5 V to 6 V, while the **RF** signal itself can take any value from 0 V to 10 V. To resolve this mismatch, the interpolation is clipped at 0 and 10 on the normalized photodiode signal axis, and this clipped boundary value is used for any signal outside that range.

Using this interpolation, the probe beam pulse shape was measured with an **SPCM** for two target pulse shapes, a Gaussian pulse and a Tukey pulse. This is shown in figure 2.11.



(a) Calibrated Gaussian pulse detected on the **SPCM**. The value set in the sequence is $\sigma = 0.5 \mu\text{s}$.

(b) Calibrated Tukey pulse detected on the **SPCM**. The values set in the sequence are the rise time $\tau_r = 2 \mu\text{s}$ and the up time $\tau_u = 3 \mu\text{s}$.

Figure 2.11: Gaussian and Tukey pulses recorded using the **SPCM** as well as the calibration described above. The measured photon rate is shown as blue circles, and the model function from which the given parameters are extracted as a black line.

Both pulses closely match their target shape, with only small deviations in amplitude and width and a slight discrepancy near the peak of the Gaussian pulse. This deviation is likely caused by the **AOM** calibration curve not transferring accurately to the pulses used here: the calibration was recorded using a comparatively slow $200 \mu\text{s}$ linear **RF** ramp, whereas the **RF** voltage during the Gaussian pulse ($\sigma \approx 0.5 \mu\text{s}$) varies on timescales roughly two orders of magnitude faster. If the **AOM**'s response is not fully rate-independent, applying a calibration curve recorded under such quasi-static conditions to much faster signals could introduce a systematic, shape-dependent distortion in the output pulse. This close

agreement is nonetheless itself evidence that the calibration works as intended: had it not, the measured pulses would be visibly distorted from the target shape, rather than showing only this small quantitative deviation.

It was further verified that only a negligible amount of residual probe light is emitted when the AOM is switched off, with an extinction ratio of $3.8 \cdot 10^4$ being found – a typical order of magnitude for AOM extinction [29]. Preventing a residual photon rate outside the intended pulse window is important, as it can otherwise affect the coherence time measurement described in chapter 4.

Overall, aside from these small discrepancies in the extracted pulse parameters, the calibration still meets the requirements of the experiment, and well-shaped Gaussian and Tukey pulses can be generated with it. A key advantage of this calibration is that, compared to model functions such as the previously used 18th-order polynomial, this interpolation behaves significantly better at the extremes of the calibration interval, since a high-order polynomial cannot follow the steep slope at the edges of the calibration curve and tends to oscillate there. Together with the extinction ratio of $3.8 \cdot 10^4$, this ensures well-defined probe pulses for the coherence time measurements.

2.3.5 Flat-Top Control Beam

To further prepare the experiment for the following coherence time and correlation measurements, an upgrade to the control beam path was made. As was mentioned in section 2.2, the control beam excites atoms from the intermediate state to the Rydberg state. Due to the weak coupling between both states, the control field has to be relatively strong. While the control field is resonant with the intermediate-to-Rydberg transition, the atoms predominantly reside in the ground state. From the ground state's point of view, the control beam is therefore simply a far off-resonant field, and any laser with a spatially varying intensity that is off-resonant from the atoms it illuminates exerts a dipole force on them. This conservative dipole force is given in [30] as

$$F_{\text{dip}}(\vec{r}) = -\nabla U_{\text{dip}}(\vec{r}) = \frac{1}{2\epsilon_0 c} \text{Re}(\alpha) \nabla I(\vec{r}). \quad (2.8)$$

Here, α is the complex polarizability and $I(\vec{r})$ is the position-dependent intensity of the control beam. Since the control light at 395 nm is blue-detuned with respect to the 399 nm ground-state transition, the resulting potential is repulsive. Due to this dipole force, atoms are pushed out of the trap; since the control beam is switched on and off over many experimental pulses, this also leads to heating. Since the force is proportional to the intensity gradient $\nabla I(\vec{r})$, a beam with a spatially uniform intensity profile exerts essentially no dipole force, even though the atoms still experience the associated uniform Stark shift.

The solution implemented here is therefore to reshape the control beam from its previous Gaussian intensity profile into a flat-top profile, so that the region occupied by the trapped atoms sees a near-constant intensity. This is done using a beam shaper⁴, which uses a diffractive optical element to change the phase of the Gaussian beam, producing a round flat-top intensity profile at the focus; since this profile is only maintained within a limited region of the beam around the focal plane, the intensity is near-constant only in this region.

⁴ FBSR-405-1.5-A from TOPAG Lasertechnik

With the newly implemented flat-top beam shape, the dipole force exerted on the atoms at the probe beam focus and in the center of the trap is suppressed, mitigating the atom loss across pulses. This change provides improved conditions for the coherence time and correlation measurements discussed in chapters 4 and 5.

2.4 Imaging of the Atomic Cloud and Detection of Photons and Ions

The measurements presented in this thesis rely on three kinds of information:

1. the position, the atom number, the atom number density, and the temperature of the atomic cloud, obtained by absorption imaging,
2. the probe photons transmitted through the atomic cloud, detected in a [Hanbury Brown Twiss \(HBT\)](#) setup, and
3. the ions created by ionization of the Rydberg atoms, detected with an [MCP](#) detector.

For absorption imaging, a beam of wavelength 399 nm and diameter 2 mm is sent through the atom cloud – from bottom to top for vertical imaging, from side to side for horizontal imaging. What remains of the beam after traversing the atoms is magnified by a factor of 1.6 and 2.5, respectively, and recorded by a camera⁵. A schematic of the imaging setup is shown in figure 2.12.




Figure 2.12: Simplified schematic of the absorption imaging setup (side view). The atomic cloud is imaged along the horizontal and the vertical axis with resonant light at 399 nm, magnified by telescopes with a magnification of 2.5 and 1.6, respectively, and recorded on a camera. The inset shows the atomic cloud and the probe beam, which propagates perpendicular to the image plane.

From the resulting absorption image, the position of the cloud and the atom number are obtained by Gaussian functions adapted to its projections [31], and the temperature from the ballistic expansion of the cloud after its release from the dipole trap [32]. The peak atomic density is then calculated from

⁵ pco.pixelfly usb

the atom number and the temperature in the harmonic approximation of the dipole trap, with the trap frequencies obtained from the dipole trap power [30].

The transmitted photons are detected using a **HB**T setup consisting of non-polarizing 50:50 beam splitters, which is shown in figure 2.13.

The probe light coupled into the collection fiber shown in figure 2.2 is guided to the **HB**T setup, where residual control light is filtered out using an interference filter. The probe light is first split by one 50:50 beam splitter, and each of the resulting two beams is split again by a further 50:50 beam splitter, yielding four outputs. These are coupled into four different multimode fibers and detected by four **SPCMs**⁶, which are connected to a TimeTagger⁷ registering each photon detection event, or ‘click’. These detectors can be used not only to measure two-photon correlations, as described in detail in chapter 5, but also to measure the overall photon counts, from which the transmission and the **OD** are obtained. Splitting the light across four detectors instead of two also reduces the fraction of coincident photon pairs lost to detector dead time, and provides multiple independent cross-detector pairs that can be compared.




Figure 2.13: **HB**T setup for detecting photons and photon correlations.

The Rydberg state population, as required, e.g., for the lifetime measurement in chapter 4, is detected via field ionization. For the Rydberg state detection, eight in-vacuum electrodes, arranged as ring segments above and below the atomic cloud, are used together with an **MCP** detector. These electrodes serve two purposes: compensating stray electric fields at the atomic cloud position, and generating the pulsed electric ionization field used to ionize the Rydberg atoms in the cloud and guide the resulting ions to the detector. The stray fields are compensated along the three coordinate axes, reaching residual fields at or below 1 mV cm^{-1} . To detect a Rydberg excitation, the atoms are field-ionized using a pulsed electric field applied via two of the eight electrodes, and the generated ions are guided onto the active area of the **MCP** detector.

Together, the three detection methods provide complementary access to the system: absorption imaging characterizes the atomic cloud and is used for the alignment of the dipole trap, the **HB**T setup records the transmission and the photon statistics of the probe light, and the ion detection probes the Rydberg

⁶ Laser Components COUNT-100B-FC

⁷ TimeTagger 20 from Swabian Instruments

state population. These form the basis of the OD measurements in chapter 3, the lifetime and coherence time measurements in chapter 4, and the correlation measurements in chapter 5.

Light Propagation in Dense Atomic Media

The OD is a key figure contributing to the strength of the photon–photon interaction mediated by the atomic medium. The Rydberg blockade breaks the EIT condition within a blockade radius r_b around an existing polariton. A photon–photon interaction becomes effective once this blockade radius exceeds the mean free path a photon travels before being scattered in a resonant two-level medium. This mean free path scales inversely with the OD of the medium, so that a strong photon–photon interaction requires a high OD [9]. This light–atom coupling is imperative for introducing interaction between photons in the medium. Thus, accurately determining the OD is very important. So far, both in this group and in the literature, the OD has been extracted from the probe transmission measured as a function of detuning. This extraction relies on three assumptions: a weak probe beam, a probe beam shape that is not changed by the atomic medium, and an atomic density that is transversally homogeneous over the extent of the probe beam.

In this chapter, it is shown that the latter two assumptions do not hold anymore at the atomic densities and cloud shapes which can be prepared in the Yb-174 apparatus. Instead, a numerical method, the ASM, is used to simulate the probe beam (see section 2.3) propagation through the dense atomic cloud as well as the resulting transmission and OD. This simulation is done for a two-level atom and is used to examine how the probe beam shape changes when propagating through the atomic cloud. From the transmitted probe light, the OD of the atomic cloud is extracted and compared to measurement results. To that end, section 3.1 revisits the standard definition of the OD and shows why, at the atomic densities reached in the Yb-174 apparatus, the assumption that the probe beam shape is unaffected by the cloud no longer holds. Section 3.2 derives the ASM used throughout this chapter to propagate the probe field through a spatially varying refractive index, and section 3.3 validates this implementation against two analytically solvable limits, the thin-lens approximation and the Lambert–Beer law. Section 3.4 then applies the method to realistic atomic cloud configurations to build an intuition of how the probe field shape changes after traversing the atoms, as well as of the influence of the cloud size and position. Further, the density of the atomic cloud is varied, and the resulting OD is compared against the OD extracted from the experiment using two independent methods. As a last step, to validate the introduced beam propagation, a measurement result from the neighboring research group Rubidium Quantum Optics (RQO) is shown.

3.1 Optical Depth in Atomic Media

In this section, the definition of the **OD** used so far in this group and in the literature for comparable experiments is introduced, in order to highlight the limitations of the model used to date. So far, the **OD** has been defined in the weak-probe limit for a closed two-level system, treating the probe as propagating along a single line through a cold atomic cloud. It is generally defined via the transmission, as shown in equations 2.5 and 2.6; these equations do not yet capture the dependence of the cross section on the probe detuning, which is introduced next.

Scanning the probe beam frequency around the resonance of $|e\rangle$ introduces an explicit detuning dependence into the cross section (see equation 2.6), giving rise to a transmission dip with a Lorentzian absorption profile, described by

$$T(\Delta_p) = \exp\left(-\frac{\text{OD}_0}{1 + \tilde{\Delta}_p^2}\right), \quad (3.1)$$

where $\tilde{\Delta}_p = \Delta_p/\gamma_{ge}$ is the normalized probe detuning and OD_0 is the peak (resonant, $\Delta_p = 0$) **OD** from equation 2.6. This Lorentzian shape arises directly from the natural linewidth of the probe transition, which alone sets the width of the dip for a cold, closed two-level atom. If the local **OD** can be treated as approximately constant over the region sampled by the beam, the position dependence of the **OD** drops out entirely and only the detuning-dependent exponential on the right remains, with **OD** now a single number rather than a function of position. This is how the **OD** is generally extracted in measurements: as a single scalar number extracted with a single model function, e.g., in [9, 33, 34].

This equation is derived for a purely one-dimensional geometry: it describes the probe as propagating along a single line, with no transverse coordinate at all. It therefore cannot capture how a real probe beam, which has finite transverse extent, is reshaped by refraction and diffraction as it propagates through a cloud whose density varies transversally. The beam experiences different refraction and diffraction conditions at different positions, reshaping it as it propagates; and since this reshaping depends on how far the probe is detuned, the cloud does not just distort the beam, it distorts it differently at each detuning.

In summary, this section has shown that while the **OD** is straightforward to define and measure via a detuning scan under the weak-probe and two-level assumptions made above, this one-dimensional treatment cannot capture the transverse atomic density variation of the cloud and the beam-reshaping effects that become significant for the cloud densities considered here. This is why a beam-propagation treatment that does not rely on it is developed in the remainder of this chapter.

3.2 Angular Spectrum Method

In this section, the **ASM** in the paraxial approximation, the numerical tool used throughout this chapter to propagate the probe field through the atomic cloud's spatially varying refractive index, is introduced.

The **ASM** can be derived from the Maxwell equations in a medium with a position-dependent permittivity $\epsilon = \epsilon(x, y, z)$, no free charges and currents, as well as a constant permeability, $\mu = \mu_0$. This derivation is based on [35, 36]. Here, $\vec{E} = \vec{E}(x, y, z, t)$ denotes the electric field vector. The resulting wave equation is then given by

$$\nabla^2 \vec{E} - \mu\epsilon \frac{\partial^2 \vec{E}}{\partial t^2} = -\nabla \left(\vec{E} \cdot \frac{\nabla \epsilon}{\epsilon} \right). \quad (3.2)$$

Under the assumption that the variation of the refractive index is small over the distance of one optical wavelength, $|\nabla \epsilon|/\epsilon \ll k_0$, and since in this approximation the electric field components are decoupled, the electric field is renamed ψ . Under the assumption that ψ predominantly propagates along the x -direction with a complex envelope ψ_e , ψ can be rewritten as

$$\psi(x, y, z, t) = \psi_e(x, y, z) \exp(-i(\omega_0 t - k_0 x)). \quad (3.3)$$

Here, ω_0 is the beam's carrier angular frequency and $k_0 = \omega_0/c$ its vacuum wavenumber. Using all the above assumptions and changes, the wave equation becomes

$$\nabla_\perp^2 \psi_e + \frac{\partial^2}{\partial x^2} \psi_e + 2ik_0 \frac{\partial}{\partial x} \psi_e - (k_0^2 - \mu_0 \epsilon \omega_0^2) \psi_e = 0. \quad (3.4)$$

Another assumption can be made, namely that ψ_e varies slowly as a function of x , the slowly varying envelope approximation, which results in $\left| \frac{\partial \psi_e}{\partial x} \right| \ll |k_0 \psi_e|$, allowing the $\partial^2 \psi_e / \partial x^2$ term above to be dropped. Further, with the common relations $k_0 = \omega_0 \sqrt{\mu_0 \epsilon_0}$ and $n = \sqrt{\epsilon_r}$, the (in general complex) refractive index, where $\epsilon_r = \epsilon/\epsilon_0$, the wave equation can be rewritten as

$$\frac{\partial \psi_e}{\partial x} = \frac{i}{2k_0} \nabla_\perp^2 \psi_e + \frac{ik_0}{2} (n^2 - 1) \psi_e. \quad (3.5)$$

For numerically propagating the beam, it is rewritten in terms of two operators, $\hat{D} = \frac{i}{2k_0} \nabla_\perp^2$ and $\hat{S} = \frac{ik_0}{2} (n^2 - 1)$, as

$$\frac{\partial \psi_e}{\partial x} = (\hat{D} + \hat{S}) \psi_e. \quad (3.6)$$

Here, $\hat{D}$ is the diffraction operator, describing the effect of propagation accounting for diffraction, and $\hat{S}$ is the inhomogeneous operator describing inhomogeneities of the medium. By utilizing the Baker–Campbell–Hausdorff formula and truncating the expansion to first order in Δx , the solution to this differential equation is given by

$$\psi_e(x + \Delta x, y, z) = \exp \left(\frac{ik_0}{2} (n^2 - 1) \Delta x \right) \mathcal{F}^{-1} \left\{ \exp \left(-\frac{i(k_y^2 + k_z^2) \Delta x}{2k_0} \right) \mathcal{F} \{ \psi_e(x, y, z) \} \right\}. \quad (3.7)$$

Here, $\mathcal{F}\{\cdot\}$ denotes the two-dimensional Fourier transform over the transverse plane (y, z) , and k_y, k_z are the corresponding transverse spatial frequencies. This equation describes the step-wise propagation of the input electric field complex envelope through the medium. For a proper simulation, this step by Δx has to be repeated for a number of defined x increments.

The Fourier-domain phase factor acts on the transverse Fourier transform of ψ_e , advancing each transverse spatial frequency component by the paraxial free-space phase it would accumulate over the step – exactly as in vacuum propagation. The real-space prefactor acts pointwise in (y, z) : since $n = n' + in''$ is

complex and position-dependent, its imaginary part (from n'') attenuates the field, capturing absorption, while its real part (from n' , the dispersive part of the refractive index) imprints a position-dependent phase shift, capturing refraction. Repeated evaluation of equation 3.7 in steps of Δx thus propagates the probe field through an arbitrary refractive index distribution $n(x, y, z)$, capturing both the beam reshaping and the absorption it undergoes along the way.

3.3 Validation of the Angular Spectrum Method

Before applying this simulation to beam propagation through unknown refractive index distributions, it is validated in this section by recreating two well-known results: the thin-lens approximation (section 3.3.1) and the Lambert–Beer law (section 3.3.2).

3.3.1 Comparison with the Thin-Lens Approximation

In this section, the *ASM*'s numerical accuracy is checked against a case with a known closed-form solution: a Gaussian beam refocused by a thin lens. Since the resulting waist and focus position can be calculated analytically, this provides a direct, quantitative benchmark of the simulation's accuracy.

A thin lens acts on an incident field by imposing a position-dependent phase shift on its wavefront. Mathematically, this corresponds to a multiplication of the field by the transmittance function [37]

$$\exp\left(-\frac{i\pi r^2}{\lambda f}\right), \quad (3.8)$$

where f is the focal length of the lens and r^2 is the squared transverse distance from the optical axis.

As a test case, a Gaussian beam is propagated through a thin lens placed at the beam's own waist position. The initial Gaussian beam has a waist of $50\ \mu\text{m}$ and a wavelength of $398.91\ \text{nm}$. When this Gaussian beam is propagated through a lens with a focal length of $5\ \text{mm}$, the resulting field is again a Gaussian beam, whose new waist and focus position can be calculated analytically. Taking the lens position as the origin, the new focus position x'_0 is given by [38]

$$x'_0 = \frac{f}{1 + (f/z_R)^2}, \quad (3.9)$$

where $z_R = \pi w_0^2/\lambda$ is the Rayleigh range of the incoming Gaussian beam, and the corresponding new beam waist by

$$w'_0 = \frac{f}{\sqrt{f^2 + z_R^2}} w_0. \quad (3.10)$$

Using these expressions, the results of the numerical simulation can be compared to the analytical predictions.

In the simulation, the complex field envelope of the Gaussian beam is propagated stepwise along x , using a step size of approximately $18\ \mu\text{m}$, first over $5\ \text{mm}$ of free propagation up to the lens (where the beam reaches its own natural focus), and then over a second stage of $9.39\ \text{mm}$ after the lens – symmetric

around the new focus, extending as far past it as the lens-to-new-focus distance itself – for a total propagated range of 14.39 mm, on a transverse grid of resolution $1.2\ \mu\text{m}$ and a transverse window of $600\ \mu\text{m} \times 600\ \mu\text{m}$. At each propagation step, the beam waist is extracted by fitting the transverse intensity profile to a Gaussian, using the model

$$I(y, z) = I_0 \exp\left(-2\left(\frac{r}{w}\right)^2\right), \quad r^2 = (y - y_0)^2 + (z - z_0)^2, \quad (3.11)$$

to the transverse intensity profile. The resulting set of waists $w(x)$ is then fitted with the analytical expression for the evolution of a Gaussian beam's waist (see equation 2.4) to extract the minimum waist w'_0 and its position x'_0 , independently along both transversal coordinates y and z . This can be seen in figure 3.1, with the extracted and theoretical values compared in table 3.1, which show that the simulated and analytical focus positions and beam waists agree to within 0.36% and 0.38%, respectively.

Table 3.1: Theoretical and simulated waist w'_0 and focus position x'_0 (measured relative to the lens) after the thin lens. The fitted values extracted along y and z agree with each other to within 8 nm in the waist and $2.2\ \mu\text{m}$ in the focus position, confirming that the beam stays circular through the whole system. Both the difference between the two directions and the deviation from the theoretical values exceed the statistical uncertainties, showing that the accuracy is limited by the systematic effects of the discretization.

	Theory	Simulation (y)	Simulation (z)
$w'_0\ (\mu\text{m})$	12.3070	12.260 ± 0.001	12.268 ± 0.001
$x'_0\ (\mu\text{m})$	4697.073	4680.1 ± 0.2	4682.3 ± 0.2




Figure 3.1: Waist evolution of a Gaussian beam refocused by a thin lens placed at the beam's initial waist. Comparison between the [ASM](#) simulation (blue dots) and the analytical prediction (black dashed line); the vertical black line marks the lens position. Due to the simulated data points not having an explicitly defined noise source, the statistical uncertainties quoted on w'_0 and x'_0 likely underestimate the true uncertainty, which is dominated by the systematic effects discussed in the text.

Both the analytical result and the simulation share the same modeling assumptions: the paraxial approximation, an infinitesimally thin, purely phase-shifting lens, and a scalar field. These factors thus

do not cause the difference between simulation and analytical result. A reason for the mismatch between both results can be the discretization of the numerical propagation, namely the finite longitudinal step size, the finite transverse grid resolution, and the finite transverse simulation window.

Aside from the numerical errors and the intrinsic errors of this propagation, the simulation reproduces the focusing of the Gaussian beam through a lens to within the agreement quantified above. This confirms that the [ASM](#) can be applied with confidence to the probe beam passing through the atomic cloud, where no closed-form solution exists and the simulation itself becomes the reference against which the measured probe transmission is compared.

3.3.2 Recovery of the Lambert–Beer Law

In addition to the thin-lens test, the [ASM](#) simulation is also validated against the Lambert–Beer law. While the thin-lens test verifies the free-space propagation by the diffraction operator $\hat{D}$, since the lens is applied as a phase mask, this test isolates the imaginary part n'' of the refractive index, responsible for absorption. Together, the two benchmarks separately confirm that the [ASM](#) correctly implements diffraction and absorption, before the two are combined with refraction for the atomic cloud's spatially varying, complex refractive index.

Analytically, the Lambert–Beer law describes the intensity decay inside a homogeneous medium as follows:

$$I(x) = I_0 e^{-\alpha(x-x_0)} = I_0 \exp\left(-\frac{4\pi n''}{\lambda}(x-x_0)\right), \quad (3.12)$$

where α is the absorption coefficient, n'' is the imaginary part of the refractive index $n = n' + in''$, and x_0 marks the position at which the medium begins. With this relation, the absorption coefficient can be calculated using set values of the system.


In this test, a spatially uniform imaginary part of the refractive index of $n'' = 10^{-4}$, a beam waist of $w_0 = 10 \mu\text{m}$, and a wavelength of $\lambda = 398.91 \text{ nm}$ are chosen, and the beam is propagated stepwise along x with a step size of approximately $1.25 \mu\text{m}$ out to $x = 1 \text{ mm}$, on a transverse grid with $0.8 \mu\text{m}$ resolution. The initial transversal intensity distribution of the beam is generated at its own waist at $x = x_0 = 0$, the same plane where the medium begins.

The absorption coefficient can be determined from the simulation by comparing the incoming and outgoing intensities at varying medium thicknesses. Specifically, the intensity is integrated over the full transverse simulation window at x_0 and at each subsequent propagation step, and the ratio of these integrated intensities is used to determine the absorption at each step. Using this method, the absorption coefficient can be calculated via

$$\alpha(x_i) = \frac{\ln(I_0) - \ln(I(x_i))}{x_i - x_0}. \quad (3.13)$$

The resulting comparison of the calculated and simulated intensity and absorption coefficient of a Gaussian beam is plotted in figures [3.2\(a\)](#) and [3.2\(b\)](#).

The calculated and simulated intensity decay agree, with no resolvable deviation visible at the scale used (figure [3.2\(a\)](#)). Another quantity for how well both results agree is the plot of the extracted absorption coefficient (figure [3.2\(b\)](#)). The ratio of the simulated to the calculated absorption coefficient deviates



(a) Analytical (black dashed line) and simulated (pink dots) intensity decay for a medium with $n = 1 + i \cdot 10^{-4}$.

(b) Relative deviation $\alpha/\alpha_{\text{theory}} - 1$ (green dots) of the ratio of the absorption coefficient derived from the simulation α and the analytically calculated absorption coefficient α_{theory} . This ratio should ideally be close to one; the residual deviation is caused by accumulated numerical error.

Figure 3.2: Comparison of the numerical simulation (ASM) to the analytical Lambert–Beer prediction: (a) intensity decay along the propagation direction, and (b) absorption coefficient extracted from the simulation relative to the theoretical value.

from unity by about 3×10^{-13} . This deviation is largest for the smallest propagation distances x_i : since $\alpha(x_i)$ is obtained by dividing by $x_i - x_0$, any fixed amount of numerical noise in $I(x_i)$ is amplified most strongly where this denominator is smallest, before decreasing again as x_i grows. Beyond this minimum, the noise rises again toward larger x_i as the accumulated numerical noise from the growing number of propagation steps becomes the dominant contribution. The machine precision is on the order of 10^{-16} , and since a number of Fourier transformations are used and this uncertainty stacks, an overall deviation several orders of magnitude above single-operation machine precision, accumulated over the several hundred propagation steps, is expected.

Aside from the intensity decay and the absorption coefficient, the phase of a Gaussian beam's complex field envelope evolves along the propagation axis according to the Gouy phase,

$$\varphi(x) = -\zeta(x) = -\arctan\left(\frac{x}{z_R}\right), \quad (r = 0), \quad (3.14)$$

with z_R the beam's Rayleigh range, as defined in the previous section.

Since the refractive index only causes a uniform amplitude attenuation, which is position-independent, the beam's phase evolution should be identical to the evolution in vacuum. Thus, this value can be extracted from the simulation and compared to the theoretical value. The reason only the on-axis phase is considered is that the overall amplitude of the Gaussian beam gets very small away from the beam center, making extracting the phase over the whole area more prone to noise. It can be seen in figure 3.3 that the magnitudes of the simulated and calculated Gouy phase agree to within a maximum relative deviation of 0.81%.

With this and the previous section, the ASM has been shown to reproduce the analytical Gaussian-beam




Figure 3.3: On-axis phase evolution (Gouy phase) of the Gaussian beam in the [ASM](#) (blue dots) compared to the analytically calculated Gouy phase (dashed black line), both shown as magnitudes on a logarithmic scale.

and Lambert–Beer predictions to within the accuracy expected from numerical discretization – agreement to within 0.4% in focus position and waist for the thin-lens test, a ratio deviation of order 10^{-13} for the absorption coefficient, and a maximum relative deviation of 0.81% for the on-axis Gouy phase – and is therefore used in the remainder of this chapter to propagate the probe beam through the atomic cloud.

3.4 Optical Depth for Different Atomic Cloud Configurations

Since the simulation has been shown to give reasonably precise results, the [OD](#) of the atomic cloud can now be modeled without the assumptions of an unchanged probe beam shape and a homogeneous transverse atomic cloud in the probe beam range identified in section 3.1, while the weak-probe assumption is retained. This section aims to build an intuition for how the cloud position, size, and atomic density affect the transmitted signal. Furthermore, it provides insight into the probe beam’s behavior, which is difficult to obtain analytically due to diffraction and refraction of the beam. To this end, the atomic cloud is modeled in addition to propagating the beam, with the cloud shape extracted from the potential generated by the crossed dipole trap beams.

To calculate this potential, the AC Stark shift is considered, by which an off-resonant laser field creates a conservative trapping potential for neutral atoms. The following derivation is based on [28, 30]. To second order in the driving electric field, the induced energy shift of an atomic state coupled to a set of other states k by a light field of intensity $I(\vec{r})$ and frequency ω_L is given by

$$U_{\text{trap}}(\vec{r}) = - \sum_k \frac{\omega_{0,k} |\vec{d}_k|^2}{3\hbar\epsilon_0 c (\omega_{0,k}^2 - \omega_L^2)} I(\vec{r}), \quad (3.15)$$

where $\omega_{0,k}$ and $|\vec{d}_k|^2 = \Gamma_k \frac{2J_{e,k}+1}{2J_{g,k}+1} \frac{3\pi\epsilon_0\hbar c^3}{\omega_{0,k}^3}$ are the transition frequency and squared reduced dipole matrix element of transition k , respectively, with Γ_k the decay rate and $J_{g,k}$, $J_{e,k}$ the total angular momentum quantum numbers of the ground and excited states of transition k . Substituting this into the equation above gives the potential for the atoms in the dipole trap,

$$U_{\text{trap}}(\vec{r}) = - \sum_k \frac{\pi c^2 \Gamma_k}{\omega_{0,k}^2 (\omega_{0,k}^2 - \omega_L^2)} \frac{2J_{e,k} + 1}{2J_{g,k} + 1} I(\vec{r}). \quad (3.16)$$

For a red-detuned trapping laser ($\omega_L < \omega_{0,k}$ for the dominant transitions), this creates an attractive potential well at the intensity maximum of the light field.

To find the atom number density $n_a(\vec{r})$ in this potential, thermal equilibrium at temperature T is assumed. The atom number density then follows a Maxwell–Boltzmann distribution, renormalized to give the total atom number in the trap,

$$n_a(\mathbf{r}) = N_a \frac{\exp(-U_{\text{trap}}(\mathbf{r})/k_B T)}{\int \exp(-U_{\text{trap}}(\mathbf{r}')/k_B T) d^3 r'}, \quad (3.17)$$

where N_a is the atom number in the dipole trap.


In the experimental setup, the trap is formed by two crossed, focused Gaussian laser beams intersecting at a half-angle ϕ . Because the two dipole trap beams do not interfere, owing to their orthogonal polarization, the total intensity is the sum of the individual beams. The combined intensity profile of these beams is used to generate the potential in equation 3.16, and thus the atom number density in equation 3.17. Since equation 3.16 describes the full, numerically calculated trap potential, it also captures deviations from a simple ellipsoidal cloud shape arising from the trap's crossed-beam geometry.

With the atomic cloud shape and the beam propagation method established, the probe beam shape after traversing the atomic cloud can be calculated. This is done in figures 3.4 and 3.5 for varying probe detunings from the transition, as well as for propagation without the atomic cloud. From these simulations, the OD experienced by the probe beam is obtained by comparing the total probe power before and after the atomic cloud. For this purpose, the intensity is integrated over the full transverse simulation window at the initial and at the final propagation step, yielding the powers P_{in} and P_{out} , and the OD is given by $\text{OD}_{\text{ASM}} = \ln(P_{\text{in}}/P_{\text{out}})$, in analogy to equation 2.5.

In figure 3.4(a), the probe beam is shown after propagating in free space for 200 μm , the same propagation distance as for the other simulations. The simulated transversal window is 80 $\mu\text{m} \times 80 \mu\text{m}$. In figures 3.4(b), 3.5(a), and 3.5(b), the probe beam propagating through a typical atom cloud is shown. The dipole trap beams have a power of 16.0 W and a beam waist of 42.6 μm , and cross at a half-angle of 10.12°; together with a typical atom number of $5.13 \cdot 10^5$ atoms at a temperature of 33 μK , these values determine the trap potential U_{trap} via equation 3.16 and the thermal atomic density via equation 3.17, from which the cloud waists in propagation direction, $w_{0,x} = 48.82 \mu\text{m}$, and transversally to it, $w_{0,y} = 8.59 \mu\text{m}$ and $w_{0,z} = 8.81 \mu\text{m}$, are calculated.

It can be seen that, depending on the detuning of the probe beam, the transmitted normalized intensity looks very different. This can be explained when considering the refractive index factoring into equation 3.7. The refractive index is related to the probe-detuning-dependent susceptibility of the medium via $n^2 = 1 + \chi$. The susceptibility is then given by

$$\chi = \frac{n_a \lambda \sigma_a (\Delta_p = 0)}{2\pi} \frac{i - \tilde{\Delta}_p}{1 + \tilde{\Delta}_p^2}, \quad (3.18)$$



(a) No atom cloud.

 (b) $\Delta_p = 0$ MHz.


Figure 3.4: Probe beam intensity in the observation plane: (a) without the atom cloud, and (b) after propagation through the cloud on resonance, where the cloud acts as a pure absorber.

with $\sigma_a(\Delta_p = 0)$ the resonant absorption cross section, as defined in section 2.3.3, and $\tilde{\Delta}_p$ the normalized detuning as defined in section 3.1.

As can be seen here, the atomic cloud affects the probe beam differently depending on the detuning, because the susceptibility – and thus the refractive index of the cloud – has both a real (dispersive) part, odd in $\tilde{\Delta}_p$, and an imaginary (absorptive) part, even in $\tilde{\Delta}_p$, whose relative weight depends on the probe detuning. On resonance ($\Delta_p = 0$ MHz), the real part of the susceptibility vanishes and the cloud acts as a pure absorber, so the transmitted light in figure 3.4(b) purely reflects the diffraction from the cloud’s absorptive density profile. Further removed from resonance, the real part grows relative to the imaginary part, and the cloud additionally affects the phase velocity of the wave in a spatially varying way.

Since this phase shift depends on the local atomic density, the cloud acts as a lens whose focusing or defocusing character depends on the sign of the detuning. For red detuning ($\Delta_p < 0$), the refractive index is increased in the denser regions of the cloud, producing a focusing effect; for blue detuning ($\Delta_p > 0$), the opposite holds and the cloud acts as a diverging lens. This lensing effect – not any difference in absorption, since the imaginary part of χ is identical for the two equal-magnitude, oppositely signed detunings $\Delta_p = \pm 20$ MHz shown here – accounts for the different intensity distributions in figures 3.5(a) and 3.5(b).

The observed crescent shape of the transmitted light in all three cloud-present intensity distributions can be explained by the cloud’s anisotropy in the transverse plane. This anisotropy arises because the two dipole trap beams cross at a finite angle: while each beam alone is axially symmetric about its own propagation axis, their combined intensity distribution – and hence the resulting trap potential and cloud – is not rotationally symmetric about the probe beam’s propagation axis. The lensing effect is thus spatially non-uniform, reshaping the beam in this way.


 (a) $\Delta_p = -20$ MHz.

 (b) $\Delta_p = 20$ MHz.

Figure 3.5: Probe beam intensity in the observation plane after propagation through the cloud at the two oppositely signed, equal-magnitude detunings $\Delta_p = \pm 20$ MHz, showing the opposite-sign lensing effect between red and blue detuning.

Because the on-resonant cloud only absorbs, with no lensing at all, its OD sets the baseline. At red detuning, focusing toward the densest part of the cloud partially compensates for the weaker off-resonant absorption relative to resonance, giving the second-largest OD; at blue detuning, defocusing compounds with that same weaker absorption instead of compensating for it, giving the smallest OD of the three.

As can be seen, the beam can be propagated using the ASM. The atomic cloud shape can be numerically directly determined from the intensity distribution of the crossed dipole trap beams and yields results which, while not entirely obvious a priori, match the qualitative predictions of the dispersive/absorptive susceptibility model derived above: the sign of the lensing follows the sign of Δ_p , and the on-resonance OD is largest. With this, it becomes possible to more quantitatively compare the OD for varying cloud shapes and positions.

3.4.1 Simulated and Measured Transmission Valley

In this section, the measured transmission valley created by a detuning scan of the probe field is compared to the transmission valley calculated from the one-dimensional OD calculation introduced in section 3.1. Further, as motivation for introducing the transmission calculated from the ASM, the difference in OD determined from the one-dimensional calculation and from the ASM is plotted. To this end, the measured and simulated transmission valleys are compared as a function of detuning, and the OD extracted from all three methods – measurement, the one-dimensional calculation, and the ASM – is compared. The one-dimensional OD is obtained by directly evaluating equation 2.6 along the beam axis using the atomic density profile of the trapped cloud, and inserting the result into equation 3.1 to obtain the corresponding transmission valley. The ASM instead numerically propagates the probe field through the full three-dimensional atomic cloud, as introduced in section 3.2 and validated in section 3.3; the

OD is then obtained from the ratio of the total probe power before and after the cloud, according to equation 2.5. In both calculations, the coupling of the transmitted probe light into the collection fiber is neglected, i.e., ideal fiber coupling is assumed, and all power within the simulation window is counted as transmitted.


To extract the OD experimentally, the atoms are exposed to a sequence of signal pulses of the probe field in the dipole trap (see section 2.1) at a set detuning from the intermediate state $|e\rangle$. This is followed by an equal number of reference pulses, in which the same probe beam transmission is recorded with no atoms present. Comparing the signal and reference transmission for this detuning, and repeating the full signal-and-reference sequence over a range of probe detunings, yields the transmission valley, from which the OD is extracted with the model function of equation 3.1. To obtain a transmission valley that is as symmetric as possible, the coupling into the collection fiber has to be optimized and the atomic cloud has to be well aligned to the probe beam focus. A residual asymmetry remains nonetheless, since the focusing and defocusing of the probe beam by the cloud at red and blue detuning, respectively, leads to a different transmission on both sides of the resonance.

Using this method, two transmission valleys at different atomic conditions are recorded. These atomic conditions are an atomic density of $(9.74 \pm 0.87) \times 10^{12} \text{ cm}^{-3}$ and a temperature of $(12.36 \pm 0.19) \mu\text{K}$, as well as an atomic density of $(2.51 \pm 0.23) \times 10^{13} \text{ cm}^{-3}$ and a temperature of $(32.21 \pm 0.66) \mu\text{K}$. The atom number, temperature, and peak atomic density are determined from absorption images as described in section 2.4.

Both measured transmission valleys are plotted in figure 3.6. Further, the transmission determined from the ASM is plotted in the same plot, as well as the one-dimensional calculation extracted from equations 2.5 and 3.1.

As can be seen in this figure, the measured, calculated, and ASM-simulated OD disagree with one another on resonance – the detuning at which the OD is by definition extracted in all three cases, since this is where the peak OD is reached. Among the three, the measured OD is smallest, the ASM-simulated OD lies in between, and the one-dimensional calculation gives the largest value. The mismatch between the measured and simulated OD is attributed to the fact that the probe beam is assumed to be Gaussian in the model, whereas in reality it is not: this non-Gaussian beam shape allows more of the probe light to propagate around the atomic cloud, which limits the maximum OD that can be reached experimentally. The mismatch between the one-dimensional calculation and the ASM simulation, in contrast, follows directly from the one-dimensional treatment's assumptions. As seen in figure 3.6(a), the calculated OD is 26% larger than the simulated value; for the transmission valley at higher atomic density, this discrepancy grows to roughly a factor of 3. This is exactly what is expected: the one-dimensional calculation considers only a single line along the beam propagation direction and the density scaling along it, implicitly treating the atomic cloud density as transversally constant and the probe beam as unchanged along its propagation. Both assumptions push the calculated OD on resonance upward; once the true beam and cloud shape are included, as in the ASM simulation, the OD on resonance drops accordingly.

This constitutes a systematic deviation between the OD calculated with the one-dimensional model and the OD experienced by the probe beam as obtained from the ASM, showing that the one-dimensional treatment – which neglects both the transverse cloud density variation and the resulting beam reshaping – cannot be used to state a physically meaningful OD for this apparatus. Since the blockade radius ideally exceeds the probe beam waist, a single Rydberg polariton samples the transversally varying



(a) Measured, calculated, and simulated transmission valley determined for an atomic cloud with density $(9.74 \pm 0.87) \times 10^{12} \text{ cm}^{-3}$ and temperature $(12.36 \pm 0.19) \mu\text{K}$.

(b) Measured, calculated, and simulated transmission valley determined for an atomic cloud with density $(2.51 \pm 0.23) \times 10^{13} \text{ cm}^{-3}$ and temperature $(32.21 \pm 0.66) \mu\text{K}$.

Figure 3.6: Measured probe-transmission valleys compared to the simulated (ASM) and closed-form Lorentzian predictions, for the two atomic conditions given above. The ASM simulation is shown as colored circles connected by a line, the one-dimensional calculation of equation 3.1 as a gray dashed line, and the measured transmission as black markers; the right axis gives the corresponding OD.

atomic density across the whole probe beam. The OD relevant for the Rydberg-mediated nonlinearity is therefore better approximated by the OD experienced by the probe beam, which represents an effective mean over the probe beam profile, than by the one along the beam axis.

3.4.2 Influence of Atomic Cloud Geometry

In section 3.4.1, it was shown that the OD calculated with the one-dimensional model is larger than the value obtained with the ASM, motivating a closer look at how the OD depends on the atomic cloud's shape and size. To additionally build an understanding of how the overlap between the probe beam and the atomic cloud affects the OD, this section also examines the OD's dependence on the cloud's position relative to the probe beam.

To simplify the simulation, and to make adjusting the cloud position and, in particular, its shape easier, the full derivation of the cloud shape from the dipole-trap-intensity-dependent potential is not used. Instead, as an approximation, a three-dimensional Gaussian distribution is used of the form

$$n_a(x, y, z) = \frac{N_a}{(2\pi)^{3/2} \sigma_x \sigma_y \sigma_z} \exp\left(-\frac{(x - x_0)^2}{2\sigma_x^2} - \frac{(y - y_0)^2}{2\sigma_y^2} - \frac{(z - z_0)^2}{2\sigma_z^2}\right). \quad (3.19)$$

With this atom number distribution, the cloud shape is determined by the extents $\sigma_{x,y,z}$, and the cloud position is determined by x_0, y_0, z_0 . Unless stated otherwise, the peak atomic density and the extents of this Gaussian cloud are chosen to match the cloud calculated from the trap potential in section 3.4. Based on this cloud shape, the transmission can be determined, and with equation 2.5, the OD can be extracted.

Initializing this cloud with different extents and at different positions yields the plots shown in figures 3.7 and 3.8, respectively.





Figure 3.7: Cloud shape OD dependence scan. The cloud is initialized at the probe beam focus, and the transversal extent of the cloud is adjusted using σ_y and σ_z at constant peak atomic density. For the different cloud extents, the OD is determined.

As can be seen in figure 3.7, the transversal cloud extent is scanned with the peak atomic density kept constant while σ_y and σ_z are varied, such that the total atom number increases with the cloud size. The scan starts at the lower left corner, where the cloud has the same transversal extent as in section 3.4. From there, the OD increases as σ_y and σ_z are increased, saturating at the upper right corner. In this limit of a large transverse cloud size, the atomic density becomes approximately homogeneous over the extent of the probe beam, so the assumption of a homogeneous transverse cloud underlying the one-dimensional model of section 3.1 is increasingly well satisfied. This is confirmed by comparing the saturated OD of 137.7 to the value of 138.3 obtained from the one-dimensional model for the same on-axis cloud parameters: the two agree closely, showing that the ASM simulation approaches the one-dimensional result in the homogeneous-cloud limit.

This same homogeneous-cloud limit is also the assumption underlying the model used to extract the OD experimentally from the measured transmission valley: it assumes that all detected light has experienced the same OD. Since real experimental clouds are not this large, this points to a more fundamental problem with how the OD has so far been determined experimentally: part of the detected light passes through regions of lower atomic density or around the cloud, so that this experimentally determined OD underestimates the OD experienced by light passing through the center of the cloud and can therefore not be called the OD of the system without also stating the assumptions under which it holds.

In figures 3.8(a) and 3.8(b), the position of the atomic cloud is varied, which yields a convolution of the cloud shape and the probe beam shape.

Thus, in figure 3.8(a), where the cloud is scanned transversally, a roughly Gaussian-shaped OD distribution can be seen, which agrees with expectations: the simulated OD is highest when the cloud is centered at



(a) Transversal cloud position OD dependence scan. The cloud is initialized with the same widths as in the previous section, with the position of initialization y_0 and z_0 changing. For the different cloud positions in relation to the probe beam focus, centered at $(0, 0)$, the OD is determined.

(b) Longitudinal and transversal cloud position OD dependence scan. The cloud is initialized with the same widths as in the previous section, with the position of initialization x_0 and z_0 changing. For the different cloud positions in relation to the probe beam focus, centered at $(0, 0)$, the OD is determined.

Figure 3.8: Position OD scan of the atomic cloud.

the probe beam focus, since at this point the overlap between both is highest.

In figure 3.8(b), the cloud position is scanned in one transversal and one longitudinal direction, again yielding the convolution of cloud and beam shapes discussed above. This scan was performed experimentally in section 2.3.3, and the simulation shows a similar OD distribution: the OD is highest at the beam focus and decreases both transversally and longitudinally away from it. The longitudinal decrease is gradual, reaching roughly half its peak value around the beam's Rayleigh range, $z_R \approx 200 \mu\text{m}$ for the $5 \mu\text{m}$ probe beam waist and 399 nm wavelength used here, while the transversal decrease is markedly sharper, reflecting the much smaller transverse beam waist.

In summary, scanning the simulated cloud's transverse size shows that the OD approaches the one-dimensional result of section 3.1 in the limit of a large, homogeneous transverse cloud. Since this homogeneous-cloud assumption also underlies the experimental model used to extract the OD, and since real experimental clouds are considerably smaller, the experimentally extracted OD systematically underestimates the OD experienced by the probe beam – a limitation compounded by the fact that the SPCM does not distinguish light refracted or diffracted around the atoms from directly transmitted light, collecting all photons coupled into the fiber regardless of their origin. Scanning the cloud position, on the other hand, shows that the resulting OD pattern is a convolution of the cloud and probe beam shapes rather than either shape alone: transversally, the OD peaks when cloud and beam are centered on each other, as expected from their overlap, and longitudinally, the same behavior – also observed experimentally in section 2.3.3 – is recovered, with the OD decaying over the beam's Rayleigh range.

3.4.3 Influence of Atomic Conditions

The goal of this section is to compare the **OD** found using the **ASM** to the **OD** determined experimentally using two methods: the transmission-valley method (see section 3.4.1), and directly comparing the probe beam intensity measured by the **SPCMs** with and without an atomic cloud present. This comparison is done for varying atomic densities and cloud sizes, since at low densities the atomic cloud only weakly perturbs the probe beam, so the assumptions underlying both experimental extraction methods still approximately hold, whereas at higher densities the beam is increasingly diffracted and refracted by the cloud, breaking these assumptions down further; the aim is to characterize how the resulting discrepancy between the models develops as the density and cloud size increase.

The transmission-valley method, introduced in section 3.4.1, extracts the **OD** by fitting equation 3.1 to a detuning scan of the probe beam. In the second method, the **OD** is computed directly from the measured transmission rates with and without the atomic cloud present, together with a reference measurement of the background light, using the photon counts recorded by the **SPCMs**. This method is applicable as long as the transmitted rate remains measurably above the background.

Both resulting **OD** values, together with the **ASM**-simulated **OD**, are shown in figure 3.9.




Figure 3.9: **OD** dependence on the atomic conditions in the dipole trap. The **OD** was measured using two methods: from the transmission valley of a probe detuning scan (green, labeled OD (1D)) and purely from the photon rates detected by the **SPCMs** (pink, labeled I_{out}/I_{in}). The simulated **OD** (turquoise circles) was then determined using the **ASM** from the ratio of transmitted to incident power on resonance.


The **OD** in figure 3.9 is plotted on a logarithmic scale, emphasizing the difference between the simulated and measured **ODs**.

The **OD** determined from the transmission valley agrees with the **OD** directly calculated from the transmitted probe rate at low densities and strongly diverges at higher densities. The divergence between the measured **ODs** is exactly what is expected, since for low **ODs**, the atoms do not yet have such a strong effect on the probe shape – at low **OD**, the beam distortion from diffraction and refraction is only a small perturbation, so the assumption of an undistorted probe beam underlying equation 3.1 still approximately holds – and thus the two **ODs** are similar. For rising atom numbers, this changes and the measured values strongly diverge.

The OD determined with the ASM shows a factor of two to three difference to the OD determined from the transmitted probe light. This divergence is attributed to the fact that the real probe beam is not a perfect Gaussian, unlike the beam assumed both in the ASM simulation and in the derivation of equation 3.1. The photon-rate method computes the OD directly from the ratio of transmitted to incident optical power reaching the SPCMs; since the real beam carries more power outside its nominal Gaussian waist than an ideal Gaussian beam of the same width, this excess power tends to skirt the atomic cloud and still couple into the collection fiber (see section 3.4), making the measured transmission higher – and the extracted OD lower – than the simulated value.

Despite this factor of two to three, the overall scaling of the photon-rate OD with atomic density closely follows that of the ASM-simulated OD. Both experimental methods therefore fail to correctly determine the OD experienced by the probe beam – the transmission-valley method diverges strongly and does not reproduce the correct scaling, while the photon-rate method, despite being offset by a constant factor, at least captures the correct scaling with atomic density.

It is currently not possible in the YQO experiment to reliably measure the OD directly due to all the limitations mentioned above. It is generally possible to measure the OD, however, by, for example, using a camera resolving the probe beam after the atoms. Such a setup was implemented by a neighboring lab, RQO, and one of their scans is gratefully included in figure 3.10.



(a) Measured transmitted probe beam shape (provided by RQO).

(b) Simulated transmitted probe beam shape.

Figure 3.10: Simulated and measured transmitted probe beam shape recorded 50 μm after the center of the atomic cloud for a 12 MHz detuning, $\approx 20 \mu\text{K}$ temperature, and ≈ 20000 atoms. The data were provided by the neighboring RQO lab.

Figures 3.10(a) and 3.10(b) show the measured and simulated transmitted probe beam shape, respectively, 50 μm after the center of the atomic cloud, showing the same crescent-shaped diffraction from the cloud's transverse anisotropy discussed for figure 3.5. The simulation uses the atom number, temperature, atomic cloud parameters, and 12 MHz probe detuning as provided by the RQO lab for this measurement, together with the Rb-87 D2 transition's wavelength of 780.24 nm and natural linewidth of 6.065 MHz, and a probe waist of 5 μm . Unlike the simulated pattern, the measured one is visibly tilted; this difference can be attributed to aberrations such as ellipticity or astigmatism of the real probe beam, which are not

included in the simulation. The same omission also explains the difference between the simulated OD of 1.74 and the measured value of 1.23, which nonetheless agree to within a factor of ~ 1.4 . Since this measurement resolves the beam shape directly, the RQO experiment is able to extract the OD from the ratio of the total transmitted power recorded on the camera with and without atoms. This corresponds to the quantity calculated in the ASM simulation and is independent of the coupling into a collection fiber and of the fit model underlying the two extraction methods used in this section. This experimental result, even though it was not recorded by the YQO experiment, lends greater credibility to the simulated probe beam propagation presented throughout this section.

In summary, the OD cannot currently be determined correctly with the YQO experiment as it stands, with both available extraction methods showing distinct, significant deviations from the OD experienced by the probe beam. A simulation using the measured cloud parameters instead gives a more accurate estimate of the OD, which is further validated by the independent camera-based measurement from the RQO lab, agreeing to within a factor of ~ 1.4 . This simulated OD should therefore be taken as the more reliable value for this apparatus going forward.

Coherence and Lifetime of Collective Ytterbium Rydberg States

In the previous chapter, the transmission of the probe beam through the atomic cloud, as well as the resulting OD were characterized within a two-level-atom picture. Building on this, in the present chapter, the atomic level scheme is expanded to a three-level system, where the third level is a highly excited Rydberg state, as shown in section 2.2. These Rydberg atoms, due to their large dipole moments and long-range interactions [15], are a valuable tool to facilitate quantum simulation, quantum gates, single-photon sources, and precision sensing, since these same interactions are what allow a single stored excitation to mediate an effective interaction between otherwise non-interacting photons. In practice, realizing these applications also requires the resulting quantum state to survive long enough to be used: the number of coherent operations achievable on a stored collective excitation [39], and the efficiency with which it can later be retrieved as a photon [10], both degrade once the storage time approaches the collective coherence time. These effects only become possible if the state does not decay too quickly, leaving enough time for the light–atom dynamics to take place.

The available time is fundamentally bounded from above by the Rydberg lifetime τ_{Ryd} , i.e., the time a single, isolated Rydberg atom is expected to survive before decaying (section 4.1). In practice, however, this bound is rarely reached: because the stored excitation lives as a collective coherence shared across many atoms rather than a single isolated Rydberg atom, it is additionally sensitive to the atoms' thermal motion, their density, and their mutual interactions. These effects degrade the coherence time τ_c of the collective state, pulling it below the single-atom lifetime.

Thus, in this chapter, both the coherence time and the Rydberg lifetime, and the gap between them, are characterized. Further, the dominant contributor to the shortening of the coherence time compared to the Rydberg lifetime is determined. The measurement of the Rydberg lifetime is presented in section 4.1. This measured Rydberg lifetime is compared to the theoretically determined Rydberg lifetime from the PairInteraction [22, 23] software. The coherence time is characterized in section 4.2. Four dephasing mechanisms are considered to find the dominant decoherence factor:

- dephasing due to Rydberg–Rydberg interactions,
- dephasing due to ions created from ground-state atoms by the strong control laser,

- thermal dephasing, and
- collision-induced dephasing.

Collision-induced dephasing is identified as the dominant mechanism – in contrast to the thermal motional dephasing that limits comparable alkali-atom platforms. For some particular cases of low temperature and high density, it was further possible to measure a periodic coherence revival, from which a Rydberg dimer binding energy is estimated (section 4.2.3), and its underlying mechanism is explained.

4.1 Rydberg State Lifetime

This section derives the theoretical model used to determine the intrinsic lifetime of a Rydberg state, which will later be compared to the experimentally extracted lifetime. This lifetime is the quantity introduced in the previous section as τ_{Ryd} : the upper bound on how long a single, isolated Rydberg atom survives before decaying, and thus the ceiling against which the coherence time of the collective state, characterized later in this chapter, will be compared.

The finiteness of the lifetime of a Rydberg state is set by two different decay channels: spontaneous emission and **Black Body Radiation (BBR)**. Following [40], at zero temperature, an atom in the state nL can decay spontaneously to a lower-lying state $n'L'$ with a rate given by the Einstein coefficient

$$A(nL \rightarrow n'L') = \frac{\omega_{nn'}^3 e^2}{3\pi\epsilon_0 \hbar c^3} \frac{L_{\max}}{2L+1} R^2(nL \rightarrow n'L'), \quad (4.1)$$

where $\omega_{nn'} = |E_{nL} - E_{n'L'}|/\hbar$ is the transition frequency, $L_{\max} = \max(L, L')$, and $R(nL \rightarrow n'L')$ the radial dipole matrix element between the two states. The total spontaneous decay rate and zero-temperature lifetime are obtained by summing this rate over all lower-lying states,

$$\gamma_0 = \frac{1}{\tau_0} = \sum_{E_{nL} > E_{n'L'}} A(nL \rightarrow n'L'). \quad (4.2)$$

At any temperature above 0 K, the atom experiences thermal **BBR**, which can resonantly drive transitions to Rydberg states both upward and downward in energy via stimulated absorption and emission. The rate of a single such transition is given by the same Einstein coefficient as above, weighted by the mean photon occupation number of the thermal field at the transition frequency $\bar{n}_\omega = \frac{1}{\exp(\hbar\omega_{nn'}/k_B T) - 1}$. Thus, the total rate of **BBR**-induced transitions can be written as

$$\gamma_{\text{BBR}} = \frac{1}{\tau_{\text{BBR}}} = \sum_{n'} W(nL \rightarrow n'L') = \sum_{n'} A(nL \rightarrow n'L') \bar{n}_\omega. \quad (4.3)$$

Here, the expression for $A(nL \rightarrow n'L')$ is evaluated for both downward and upward transitions [40]. The Rydberg lifetime is then given by

$$\frac{1}{\tau_{\text{Ryd}}} = \frac{1}{\tau_0} + \frac{1}{\tau_{\text{BBR}}}. \quad (4.4)$$

This matches the model used by the PairInteraction software [22, 23] to calculate τ_{Ryd} for the states of interest, which is compared in section 4.1.1 to the lifetime extracted from the measurements presented later in this chapter.

In summary, this section has derived the Rydberg lifetime τ_{Ryd} from its two constituent decay channels – zero-temperature spontaneous emission and finite-temperature BBR-driven transitions – following the same model used by the PairInteraction software. With this theoretical prediction in hand, section 4.1.1 presents the corresponding experimental measurement of τ_{Ryd} and compares the two.

4.1.1 Measurement and Comparison to Theory Predictions

Having introduced, in section 4.1, the theoretical framework underlying the Rydberg lifetime, in this subsection, τ_{Ryd} is measured directly and compared to the numerical value that framework predicts as computed by PairInteraction. This comparison establishes the experimental single-atom bound that the collective coherence time discussed later in this chapter can never exceed, and, since PairInteraction is widely used to predict Rydberg-state properties, also tests how much its calculated lifetimes can be trusted.




Figure 4.1: Decay of the measured ions per pulse for varying delays between turning off the laser fields and the ionization pulse. The measured data (blue dots) are fitted with an exponential decay with a constant offset (black line), from which the Rydberg lifetime is extracted.

The lifetime can be experimentally determined by exciting the atoms in the atomic cloud in the dipole trap to the Rydberg state using the Raman (off-resonant two-photon) excitation scheme introduced in section 2.2. The probe and control lasers are then turned off simultaneously, and after a variable delay, an ionization pulse ionizes the created Rydberg atoms (not Rydberg state specific); the resulting ions are steered by the steering electrodes onto the MCP detector and counted. Repeating this for a range of delays between switching off the light fields and firing the ionization pulse yields the number of detected ions per pulse as a function of delay. With this measurement, an exponential decay of the ions over the ionization delay can be extracted. This is shown for an example measurement for the $n = 52$ Rydberg state with atom number $N_a = (5.54 \pm 0.028) \cdot 10^5$ and temperature $T = (12.75 \pm 0.10) \mu\text{K}$ in figure 4.1.

As can be seen, the extracted lifetime for this state is $(28.94 \pm 0.47) \mu\text{s}$. It should further be noted

that the exponential model does not decay to zero ions per pulse, the level expected theoretically. A possible explanation for this offset is that the ionization pulse ionizes every Rydberg state above a certain threshold, not only the initially excited state. Consequently, atoms that are transferred – for example via [BBR](#)-induced transitions – to other Rydberg states that still lie above this ionization threshold continue to be detected. If these states, e.g., higher-lying or high- L states, decay considerably slower than the initially excited state, their contribution decays only little within the measured range of delays and appears as an approximately constant offset.

The same procedure can be repeated at different Rydberg states and has been done for the states listed in table 4.1.

Table 4.1: Measured and calculated [22, 23] Rydberg state lifetimes.

n	Measured lifetime (μs)	Calculated lifetime (μs)
42	15.09 ± 0.28	12.52
52	28.94 ± 0.47	24.35
73	72.6 ± 5.1	66.73
96	353 ± 23	145.05

The measured Rydberg lifetimes can then be compared to theory. A widely used tool for determining Rydberg atom properties is the PairInteraction software, which is also used here. The calculated values for the lifetime are also shown in table 4.1. As can be seen, the measured and calculated lifetimes deviate systematically for all four Rydberg states considered, though only modestly for $n = 42, 52, 73$ and far more strongly for $n = 96$. While the detection scheme itself likely contributes to this deviation – since, as noted above, it is not specific to the initially excited state and also counts atoms transferred to nearby states – no other significant source of experimental error was identified, and the remaining discrepancy is attributed to PairInteraction’s calculation of the lifetime. It is therefore worth revisiting how the lifetime is determined (see section 4.1). The calculation of the lifetime includes the radial dipole matrix element between the initial Rydberg state and the final state reached via spontaneous emission or [BBR](#). For spontaneous emission, these final states are often low-lying states, as can be seen in figure 4.2.

The multi-channel quantum defect theory underlying PairInteraction’s calculation may produce inaccurate results for effective principal quantum numbers $\nu < 25$. The model becomes increasingly unreliable for these small ν , producing inaccurate matrix elements and energies; due to missing data, some states in this regime may not even be present in the database.¹ Under normal circumstances, this would not be a concern, since the Rydberg states considered here are, by definition, highly excited and hence have large ν . However, of the two decay channels entering the lifetime calculation, spontaneous emission also couples the Rydberg state to low-lying, low- ν states. These two effects are, as previously mentioned, [BBR](#) and spontaneous emission. While [BBR](#) mainly couples states with principal quantum numbers similar to that of the initially excited Rydberg state, spontaneous emission predominantly couples to much lower-lying states, which are strongly affected by the uncertainties of the PairInteraction calculation.

As can be seen in figure 4.2, the single largest rate is the spontaneous-emission transition from the Rydberg state to a low-lying state. This holds across all four Rydberg states considered here: the peak spontaneous-emission rate exceeds the peak [BBR](#) rate by roughly an order of magnitude at $n = 42$,

¹ Runtime output of PairInteraction, version 2.7.0.




Figure 4.2: **BBR** and spontaneous emission transition rates to other states from the Rydberg state with $n = 42$. The histogram columns show in green the transition rates generated by spontaneous emission and in purple the transition rates for **BBR** at room temperature 300 K.

narrowing to a factor of ~ 5 by $n = 96$, but **BBR** never overtakes it (see appendix A). The same holds for the summed rates: for all investigated states, the total spontaneous decay rate γ_0 exceeds the total **BBR**-induced rate γ_{BBR} at 300 K. The dominant channel setting the calculated lifetime is therefore the same low- ν spontaneous-emission transition across the full range of Rydberg states measured here. A closer look at table 4.1 shows that for $n = 42, 52, 73$, the measured lifetimes exceed the calculated ones by roughly 9%–21%. This can be explained both by the uncertainty on the calculated Rydberg lifetimes introduced above, and by **BBR** transitions to nearby Rydberg states, which are then ionized alongside the initially excited state by the same high-voltage ionization pulse. The $n = 96$ state deviates far more strongly, by a factor of 2.4 – a total gap of $\approx 208 \mu\text{s}$. Part of this gap is plausibly attributable to the Rydberg lifetime being inherently harder to measure reliably at such high n : the atoms are highly susceptible to external electric fields, and the ionization field applied during one pulse can still affect the Rydberg atoms of the subsequent pulse due to field ringing after switch-off. The remainder can likewise be explained by the same uncertainty on the spontaneous-emission contribution calculated with PairInteraction, together with the ionization of nearby Rydberg states populated via **BBR**.

In summary, the Rydberg lifetime was measured directly for all four principal quantum numbers and compared to the PairInteraction predictions. For $n = 42, 52, 73$, the measured values exceed the calculated ones by 9%–21%, consistent with two effects: the low- ν inaccuracy PairInteraction itself flags in its spontaneous-emission matrix elements – the single largest decay channel at every state considered here – and the non-specific ionization detection, which also counts atoms transferred via **BBR** to nearby states. At $n = 96$, this deviation grows sharply to a factor of 2.4, only partly attributable to the added difficulty of measuring long lifetimes at high n ; the remainder is consistent with the same two effects being more pronounced at this state, where **BBR** contributes comparatively more. Despite this discrepancy, the measured values are direct experimental results and therefore serve as the upper bound against which the collective coherence time is compared; due to the non-specific detection, this bound is to be understood as an upper estimate.

4.2 Coherence Times

In the last section, the lifetime of a single Rydberg atom, i.e., how long the population survives in $|r\rangle$ without any outside influence from other atoms, was considered. In this section, this is extended to the coherence of a stored excitation, as detailed in section 4.2.1. This stored excitation is no longer dependent on a single Rydberg state, but on ensembles of ground state and Rydberg state atoms. When traveling through the atomic cloud, the photon proceeds as a superposition of a photonic and a Rydberg component.

This duality is used to store the photon in the atomic medium as a Rydberg polariton. This is experimentally done by measuring in the EIT configuration, i.e., with the probe and control detuning $\Delta_p = \Delta_c = 0$ Hz (see figure 4.3). In this configuration, for a high control field, a Gaussian probe pulse is sent into the atom cloud. The timing is chosen such that the control field is ramped down when about half of the Gaussian pulse is stored in the atomic medium while the other half has already propagated through the atomic cloud. The stored half is stored as a stopped Rydberg polariton. To read this Rydberg polariton out again, the control field is turned on, allowing the excitation to travel through the atomic medium and to be detected by the SPCMs. The mechanisms that degrade this coherence are discussed in section 4.2.2, and the periodic coherence revival caused by Rydberg dimers is discussed in section 4.2.3.




Figure 4.3: Experimental timing to measure the storage and retrieval efficiency. The control field is shown in purple, while the probe field is shown in blue. Without the atomic cloud, in a reference measurement, the Gaussian probe pulse propagates unchanged (blue dashed line), while with atoms present, the light still in the atom cloud at the time of turning off the control pulse is stored and can be retrieved after a variable delay time, the storage time. The photons not in the medium at this time are not stored and propagate through (adapted from [41]).

4.2.1 Rydberg Collective State Coherence

As mentioned above, the coherence time measurement relies on the propagation of a polariton through the atomic medium under EIT conditions. In the following, a short overview of the storage of Rydberg polaritons in a Rydberg ensemble is given, based on [9, 14, 42, 43]. The polariton is introduced as a superposition of a photonic state and an atomic excitation, and the resulting EIT transmission window that permits its propagation through the medium is characterized. The group velocity at which the polariton propagates through the medium is then derived from the resulting dispersion. Finally, it is shown how the Rydberg lifetime sets an upper bound on the coherence time.

Storing photons in the atomic medium relies on the ladder scheme $|g\rangle \rightarrow |e\rangle \rightarrow |r\rangle$ addressed by the probe and control fields: in the steady-state solution in the weak-probe limit, the intermediate state $|e\rangle$

carries no population. The corresponding optical coherence can then be eliminated from the equations of motion, leaving a closed equation of motion for the polariton.

This concept relies on the Rydberg state's long radiative lifetime. In a ladder scheme without such a Rydberg level, this would not be possible, since the top level would decay on the same time scales as the intermediate state. Because the excitation never populates $|e\rangle$, it is not purely photonic but a joint excitation of the field and the atomic medium – a dark-state polariton [14, 43], expressed as

$$P(z, t) = \cos \vartheta \mathcal{E}_p(z, t) - \sin \vartheta \sqrt{N} \rho_{gr}(z, t) e^{i\Delta k z}, \quad \tan^2 \vartheta \approx n_g, \quad (4.5)$$

where ρ_{gr} is the coherence between the ground and Rydberg states, and n_g is the group index of the probe pulse in the atomic medium.

To determine the group index at low probe powers, the susceptibility is considered. Close to two-photon resonance, $\Delta_p = -\Delta_c$, the susceptibility of the three-level EIT medium can be approximated in terms of the two-level susceptibility introduced in equation 3.18 as

$$\chi_{3\text{-level}} \approx \chi_{2\text{-level}} \left(1 - \frac{\Gamma_{\text{EIT}}}{\Gamma_{\text{EIT}} - i\Delta_p} \right), \quad (4.6)$$

where $\Gamma_{\text{EIT}} = \Omega_c^2/4\gamma_{ge}$ is the EIT linewidth and γ_{ge} is the decoherence rate of the probe transition $|g\rangle \leftrightarrow |e\rangle$ [9]. On two-photon resonance, $\Delta_p = 0$, the group index follows as

$$n_g \approx \frac{\omega}{2} \frac{\partial \chi_R}{\partial \omega} \approx \frac{\omega}{2\Gamma_{\text{EIT}}} \text{Im}[\chi_{2\text{-level}}(\Delta_p = 0)], \quad (4.7)$$


where ω is the (angular) frequency of the probe field and $\chi_R = \text{Re}[\chi]$ is the real, dispersive part of the atomic susceptibility.

At low probe power, $n_g \propto n_a/\Omega_c^2$, i.e., for a given atomic cloud, n_g is set by the control Rabi frequency Ω_c alone. As can be seen, the character of the polariton, photonic or Rydberg, is set entirely by the mixing angle ϑ , which – through n_g – inherits the same dependence. Further, the group index sets the propagation group velocity of the polariton as $v_g = c/n_g$.

From equations 4.6 and 4.7, two experimentally measured quantities can be explained. Due to the medium response resulting from the three-level susceptibility, an EIT window around resonance appears, allowing transmission of light that is not possible in a pure two-level system. The resulting EIT window for a scan of the probe field detuning is shown in figure 4.4(a). Both this scan and the slow-light delay discussed below were measured under the same conditions: for the $n = 97$ Rydberg state, with atom number $N_a = 2.6 \times 10^5$, temperature $T = 12.3 \mu\text{K}$, and control power $P_c = 70 \text{ mW}$.

The associated group delay of the probe pulse propagating through the atomic medium under these conditions, measured by comparing a probe pulse transmitted with and without the atomic cloud present, is shown in figure 4.4(b) and was found to be $\tau_g = (1.461 \pm 0.014) \mu\text{s}$. This slow-light delay varies with the control field power.

Finally, the mixing angle between the photonic and Rydberg character of the polariton can be understood. For a strong control field ($\vartheta \rightarrow 0$), the polariton is almost entirely photonic and propagates near c ; as the control field is reduced ($\vartheta \rightarrow \pi/2$), it becomes an almost purely atomic excitation, delocalized coherently over the ensemble as a single collective Rydberg excitation,



(a) Probe detuning scan showing the full absorption of the two-level system ($|g\rangle \rightarrow |e\rangle$) on resonance (blue) and the three-level EIT window opened by the control field within the atomic absorption resonance (pink). The data points are shown together with the corresponding model curves (solid lines).

(b) Slow-light delay of the probe pulse: propagation through the atomic medium (signal, green) is delayed by $\tau_g = (1.461 \pm 0.014) \mu s$ relative to a reference measurement without atoms present (reference, purple).

Figure 4.4: EIT transmission window (left) and the corresponding slow-light delay (right), both measured for the $n = 97$ Rydberg state with $N_a = 2.6 \times 10^5$ atoms at $T = 12.3 \mu K$ and a control power of $P_c = 70$ mW.

$$|1\rangle_{\text{ph}} \otimes |g_1, \dots, g_N\rangle \longleftrightarrow |0\rangle_{\text{ph}} \otimes \frac{1}{\sqrt{N}} \sum_{j=1}^N |g_1, \dots, r_j, \dots, g_N\rangle. \quad (4.8)$$

This is the sense in which a single photon and a single Rydberg excitation coexist in superposition across the cloud: the excitation is not localized on any one atom but shared symmetrically among all N of them. Reversing the rotation of ϑ retrieves the photon. For high control powers, the polariton inherits a more photonic character, while for lower control powers, the character becomes more that of a collective excitation.

With this context, the polariton can be stored in the atomic medium by turning off the control light. As previously introduced, a single (non-collective) Rydberg excitation has a finite lifetime. Over a finite storage time t_s , this finite lifetime directly limits the polariton's retrieval efficiency, as shown in the following derivation, obtained by evolving the system under the Lindblad master equation.

Neglecting the transverse amplitude and the longitudinal phase of the stored excitation (see section 4.2.2), the collective ground state, the collective Rydberg excitation, and the collective dimer excitation are given by

$$|G\rangle = |g_1, \dots, g_N\rangle, \quad (4.9)$$

$$|R\rangle = \frac{1}{\sqrt{N_R}} \sum_{j=1}^{N_R} |g_1, \dots, r_j, \dots, g_N\rangle, \quad (4.10)$$

$$|D\rangle = \frac{1}{\sqrt{2N_D}} \sum_{j=1}^{N_D} \left(|g_1, \dots, (g_j, r_j), \dots, g_N\rangle + |g_1, \dots, (r_j, g_j), \dots, g_N\rangle \right), \quad (4.11)$$

where N_R is the number of atoms that can be excited to the Rydberg state $|r\rangle$ without forming a molecule, N_D is the number of atom pairs that can form a Rydberg dimer, and (g_j, r_j) denotes the atom pair j forming a Rydberg molecule [41]. A system with the initial state $|\psi(0)\rangle = |R\rangle$ and the Hamiltonian $\hat{H} = E_R|R\rangle\langle R| + E_D|D\rangle\langle D|$ is considered, where E_R and E_D correspond to the energies of a single Rydberg atom and a single Rydberg molecule, respectively.

The system evolves under the Lindblad master equation

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar} [\hat{H}, \rho(t)] + \sum_{n \in \{R, D\}} \frac{1}{2} \left(2\hat{L}_n \rho(t) \hat{L}_n^\dagger - \rho(t) \hat{L}_n^\dagger \hat{L}_n - \hat{L}_n^\dagger \hat{L}_n \rho(t) \right). \quad (4.12)$$

Here, the $\hat{L}_n$ are jump operators and are defined as

$$\hat{L}_R = \sqrt{\gamma_{\text{Ryd}}} |G\rangle\langle R|, \quad (4.13)$$

$$\hat{L}_D = \sqrt{\gamma_d} |G\rangle\langle D|, \quad (4.14)$$

where $\gamma_{\text{Ryd}} = 1/\tau_{\text{Ryd}}$ and γ_d are the decay rates of a single Rydberg atom and a single Rydberg molecule, respectively. Neglecting further dephasing mechanisms, the collective states decay with these single-particle rates, since each atom (each pair) decays independently with the same rate [44].

For the pure Rydberg initial state, the storage and retrieval efficiency can be extracted from the overlap of the initial and time-evolved state,

$$\eta(t_s) = \langle \psi(0) | \rho(t_s) | \psi(0) \rangle = \left| e^{-\gamma_{\text{Ryd}} t_s / 2} \right|^2 = e^{-t_s / \tau_{\text{Ryd}}}. \quad (4.15)$$

If a superposition of dimers and Rydberg excitations is stored in the atomic medium, the initial state is instead $|\psi(0)\rangle = \alpha|R\rangle + \beta|D\rangle$, with complex amplitudes α, β satisfying $|\alpha|^2 + |\beta|^2 = 1$; the resulting coherence decay is derived in section 4.2.3. Beyond this population decay, the collective coherence can be further degraded by the atoms' motion, by collisions, and by Rydberg–Rydberg interactions – the dephasing mechanisms addressed in section 4.2.2. The longitudinal dephasing and the transverse mode mismatch derived in section 4.2.2 result from the longitudinal phase and the transverse amplitude neglected above, combined with the atomic motion, and multiply directly onto the survival probability. Collisional dephasing, by contrast, is modeled as a hard-sphere collision rate γ_{coll} and is treated separately.

In summary, this subsection developed the polariton picture underlying the coherence time measurement: the propagation of a photon through the EIT medium as a joint excitation of the field and the atomic

ensemble, governed by the mixing angle ϑ and group index n_g . Tuning the control field from strong to weak rotates the polariton from an almost purely photonic excitation to an almost purely atomic one, delocalized coherently as a single collective Rydberg excitation across the ensemble. The polariton can then be stored by switching off the control field entirely. The efficiency of this storage and subsequent retrieval of the polariton is fundamentally limited by the finite Rydberg lifetime τ_{Ryd} , so that τ_{Ryd} sets the upper bound on the achievable coherence time. Beyond this population decay, the collective coherence is further degraded by motional, collisional, and interaction-induced dephasing, which are addressed in the remainder of this chapter.

4.2.2 Dephasing Mechanisms of the Rydberg Ensemble

As mentioned above, the collective coherence is further degraded by dephasing beyond the Rydberg population decay itself, and it remains to be determined which dephasing mechanism dominates the coherence time in the measurements that follow. To this end, the dephasing caused by the probe and control field powers is determined in sections 4.2.2.1 and 4.2.2.2. Further, models describing the dephasing caused by the temperature-dependent longitudinal and transversal motion of the atoms are introduced in section 4.2.2.3. Lastly, to determine the density-dependent dephasing rate, a hard-sphere collision model is introduced and the resulting dephasing is determined in section 4.2.2.4.

4.2.2.1 Rydberg–Rydberg Interaction




Figure 4.5: Dependence of the coherence time on the probe power and the resulting Rydberg–Rydberg interaction for the Rydberg states $n = 42$ (green) and $n = 52$ (purple). For larger probe power, the Rydberg–Rydberg interaction increases and the coherence time τ_c decreases.

The dephasing caused by Rydberg–Rydberg interaction can be increased by increasing the probe power. This increased probe power at the same time increases the number of created Rydberg polaritons. With the increase of the number of Rydberg atoms, the probability of two of those atoms interacting in turn increases. This interaction can shift the two-photon resonance condition of the Rydberg atoms via a van der Waals shift affecting the Rydberg level. This effectively detunes the Rydberg atom from the conditions its coherence depends on, causing decoherence.

This can be measured using the storage and retrieval efficiency measurement sequence introduced above. The same measurement is then repeated for varying Rydberg states and probe powers. The extracted coherence times from this measurement can be seen in figure 4.5. As can be seen in this figure, the coherence time decreases for larger probe powers, exactly as expected for both Rydberg states.

This effect scales with the principal quantum number, being more pronounced for higher Rydberg states, since the Rydberg blockade radius grows with n (see section 2.2): at fixed atom density and probe photon rate, the mean distance between the stored excitations is set by the number of stored photons, so that a larger blockade radius corresponds to a stronger interaction at this distance, increasing the interaction shift and hence the dephasing effect described above. Thus, to suppress the dephasing caused by too high probe powers, the experiment is run in the regime where the dephasing effect is not too strong. At the same time, the probe power cannot be chosen too low, since too few photons are then stored and retrieved per measurement cycle to extract a reliable coherence time. The probe photon rate was therefore chosen as a compromise between these two limits, on the order of $0.5 \mu\text{s}^{-1}$.

4.2.2.2 Ion Creation

Compared to Rb-87, which is commonly used in Rydberg quantum optics, Yb-174 features a couple of key differences. One of those is that the control and probe fields are very close in wavelength. In the counter-propagating beam configuration used here, this greatly reduces the recoil momentum imparted to the atoms by the two beams, but it also introduces a disadvantage that was not previously considered. As shown in section 2.2, although the probe and control light are close in wavelength – and hence in photon energy – the shorter-wavelength control light carries the higher energy. Since the control field also couples to the weaker transition $|e\rangle \leftrightarrow |r\rangle$, it has to be driven at high power. This combination sets up exactly the problem that can occur: a ground state atom can be ionized in a process called two-photon ionization. This nonlinear process runs by a ground state atom absorbing two control field photons simultaneously, resulting in the atom ejecting an electron [25].

Thus, to test the effect of the control power on the atoms, the coherence time was extracted from storage and retrieval efficiency measurements taken at varying control powers. For the considered Rydberg states, the extracted coherence times agreed within their uncertainties across the tested range of control powers from 50 to 250 mW, and no systematic trend was observed at any of the considered Rydberg states ($n = 42, 52, 73, 96$).

Ion creation by the control field is therefore excluded as a significant decoherence mechanism for the storage times and control powers used in this work. This can be attributed to the control field being switched off during the storage time, such that ions can only be created during the comparatively short write and read pulses. This is consistent with the measurements in chapter 5, where no ionization-induced structure is observed for short probe pulses, while it appears for long probe pulses, during which the control field is continuously switched on.

4.2.2.3 Longitudinal Dephasing and Transverse Mode Overlap

In alkali atoms such as Rb-87, which are typically used for Rydberg-EIT experiments, a dominant dephasing factor is the thermal dephasing caused by atomic motion, especially along the propagation direction of the excitation fields. In a counter-propagating setup of the excitation lasers such as the one used here, the atoms in the atomic cloud experience an overall momentum transfer, which increases

with a mismatch in wavelength of the excitation fields, and correspondingly limits the coherence time – for [Rb-87](#), where the probe and control wavelengths are widely separated, to only a few microseconds. Since the excitation wavelengths of [Yb-174](#) lie close together, this thermal dephasing was thought to be suppressed, promising a substantially longer thermal coherence time – one of the motivations for using [Yb-174](#) for this measurement in the first place. Whether this suppression indeed holds is checked in this section.

Thus, in this section, two thermal mechanisms that limit the storage and retrieval efficiency are introduced: the longitudinal dephasing and the transverse mode overlap. The longitudinal dephasing is caused by atomic motion along the probe beam's propagation direction, scrambling the phase of the collective excitation. The transverse mode overlap decreases due to motion perpendicular to it, which redistributes the excited atoms relative to the probe mode, such that the stored excitation no longer matches it.

As shown below (equation [4.16](#)), the distinction between both mechanisms follows directly from two contributions to the wavefunction of the stored collective excitation (with the control field turned off),

$$|R\rangle = \sum_j^N c_W(y_j, z_j) e^{i\Delta k \cdot x_j} |g, \dots, r_j, \dots, g\rangle, \quad (4.16)$$

where $c_W(y_j, z_j)$ is the amplitude envelope set by the transverse field profile of the probe beam and the exponential describes the phase evolution. The sum runs over all N atoms in the ensemble; x_j is the position of atom j along the probe propagation (longitudinal) direction, while y_j and z_j are its transverse coordinates. Δk is the wavevector mismatch between the probe and control light fields, and $|g, \dots, r_j, \dots, g\rangle$ denotes the collective excitation in which atom j alone is transferred to the Rydberg state r , while the remaining $N - 1$ atoms stay in the ground state. As can be seen from equation [4.16](#) and will be shown in the following, the amplitude envelope c_W gives rise to the transverse mode overlap, while the phase evolution $e^{i\Delta k \cdot x_j}$ gives rise to the longitudinal dephasing, so the two mechanisms can be treated independently. Due to the assumed transverse symmetry of the probe beam, the transverse mode overlap is first calculated explicitly for a single transverse coordinate, y .

First, the longitudinal dephasing is considered. It is derived from equation [4.16](#) by comparing the overlap of the polariton wavefunction at storage times $t = 0$ and $t = t_s$, with each atom's position evolving as $x_j(t_s) = x_j(0) + v_j t_s$, where v_j is the atom's velocity along x [[45](#)],

$$\begin{aligned} \eta_{\text{long}}(t_s) &= |\langle R(t=0) | R(t_s) \rangle|_{\text{long}}^2 = \left| \frac{1}{N} \sum_j \exp(i\Delta k \cdot v_j t_s) \right|^2 \\ &\xrightarrow{N \rightarrow \infty} \left| \int f(v) \exp(i\Delta k \cdot v t_s) dv \right|^2 = \exp\left(-\frac{t_s^2}{\tau_{\text{long}}^2}\right). \end{aligned} \quad (4.17)$$

Here, $f(v)$ is the thermal velocity distribution of the atoms, and τ_{long} is the coherence time limited by the longitudinal atomic movement. In this equation, a large number of atoms N is assumed, allowing the sum to be replaced by a continuous integral. The thermal velocity distribution is given here by the normalized Maxwell–Boltzmann distribution in one dimension,

$$f(v) = \frac{1}{\sqrt{2\pi} \sigma_v} \exp\left(-\frac{v^2}{2\sigma_v^2}\right), \quad \sigma_v^2 = \frac{k_B T}{m}, \quad (4.18)$$

with m the atomic mass and σ_v the one-dimensional rms velocity. Thus, the dephasing time constant is given by $\tau_{\text{long}} = (\Delta k \cdot \sigma_v)^{-1}$.

The transverse mode overlap can be determined by estimating the effect of the atoms' thermal motion on the overlap between the stored excitation and the probe mode after a storage time t_s . Upon retrieval, the amplitude $c_W(y_j(0))$ that each atom carries from its initial position is compared to the value $c_W(y_j(t_s))$ at the position the atom has moved to during the storage time. Since only the envelope carries the transverse coordinate dependence, the exponential phase factor in $|R\rangle$ cancels, and the resulting overlap is given by [46]

$$\eta_{\text{trans}}(t_s) = \left| \sum_j c_W^*(y_j(0), z_j(0)) c_W(y_j(t_s), z_j(t_s)) \right|^2, \quad (4.19)$$

where the amplitudes are normalized as $\sum_j |c_W(y_j, z_j)|^2 = 1$, and this norm of the stored excitation is conserved during storage. As before, only the transverse coordinate y is treated explicitly here; the full transverse overlap follows by symmetry and is calculated later. Further, $y_j(t_s) = y_j(0) + v_{y,j} t_s$ describes the atom's motion over time. Over the extent of the probe beam close to or at the focus, the atomic density is treated as constant (homogeneous cloud). Further, collisions are neglected here, since they are treated separately in section 4.2.2.4, so each atom moves along a straight trajectory at a constant velocity v .

The initial phase-space density of the atoms at $t = 0$ carrying the stored excitation is proportional to the product of the atom cloud's thermal velocity distribution and the spatial profile of the excitation probability set by the probe beam shape, and can be expressed as [32]

$$n_v(t_s, y, v) = f(v) \exp\left(-\frac{(y - vt_s)^2}{w_0^2}\right). \quad (4.20)$$

Here, the first factor is the one-dimensional Maxwell–Boltzmann velocity distribution of equation 4.18 at temperature T , and the second is the Gaussian beam transverse profile at the initial position of the atoms, with w_0 the waist of the probe beam. In this calculation, it is assumed that the velocity distribution is the same for all atoms, regardless of whether they are excited or not. Starting from this equation, this amplitude n_v is integrated over the velocity v to obtain the velocity-averaged amplitude of the stored excitation,

$$n(t_s, y) = \int_{-\infty}^{\infty} n_v(t_s, y, v) dv = \frac{w_0}{W} \exp\left(-\frac{y^2}{W^2}\right), \quad W^2 = w_0^2 + 2\sigma_v^2 t_s^2. \quad (4.21)$$

The overlap between this amplitude at $t = 0$ and a variable storage time t_s can then be determined using

$$\eta_{\text{trans,1d}}(t_s) = \frac{\left| \int n(t=0, y) \cdot n(t_s, y) dy \right|^2}{\left(\int |n(t=0, y)|^2 dy \right)^2}. \quad (4.22)$$

Since the stored excitation is conserved during storage, the overlap is normalized to its initial norm [46]. Evaluating this integral yields the one-dimensional transverse mode overlap,

$$\eta_{\text{trans,1d}}(t_s) = \frac{w_0^2}{w_0^2 + \sigma_v^2 t_s^2}. \quad (4.23)$$

Assuming the motion in the transverse plane is symmetric (a valid assumption, given that the probe beam shape is assumed to be radially symmetric), the overlap after a time t_s in the transverse plane can be calculated by squaring the one-dimensional overlap: $\eta_{\text{trans}}(t_s) = \eta_{\text{trans,1d}}(t_s)^2$.

Thus, the theory behind both thermal mechanisms is established. Since the longitudinal dephasing, the transverse mode mismatch, and the Rydberg-lifetime-limited decay are independent processes acting simultaneously on the same stored coherence, their combined effect is given by the product of the individual efficiency functions,

$$\eta_{\text{tot}}(t_s) = \eta_{\text{long}}(t_s) \cdot \eta_{\text{trans}}(t_s) \cdot \eta_{\text{Ryd}}(t_s) = \exp\left(-\frac{t_s^2}{\tau_{\text{long}}^2}\right) \left(\frac{w_0^2}{w_0^2 + \sigma_v^2 t_s^2}\right)^2 e^{-t_s/\tau_{\text{Ryd}}}. \quad (4.24)$$

These predictions can now be compared to the measured coherence times. However, the coherence time varies significantly with atomic density and temperature, and keeping the atomic conditions constant across measurement campaigns spanning several weeks proved difficult. Nonetheless, the measured coherence times for the different Rydberg states for similar atomic conditions, at which density-driven effects should have only a small influence, are listed in table 4.2. These values were chosen as those closest to peak atomic densities of $4.5 \times 10^{12} \text{ cm}^{-3}$ and temperatures of 7 μK ; for $n = 73$ and $n = 96$, no measurements under fully comparable conditions are available.

Table 4.2: Measured coherence times τ for each Rydberg state n , selected from the measurement points near $4.5 \times 10^{12} \text{ cm}^{-3}$ and 7 μK with the smallest uncertainty on τ .

n	ρ (10^{12} cm^{-3})	T (μK)	τ (μs)
42	4.29 ± 0.77	6.77 ± 0.20	10.90 ± 0.20
52	4.50 ± 1.97	6.34 ± 1.65	17.50 ± 0.40
73	1.00 ± 1.06	9.68 ± 0.88	16.92 ± 0.65
96	3.10 ± 1.12	11.08 ± 2.22	10.55 ± 0.26

Together with the theoretical results for the longitudinal dephasing and the transverse mode overlap, the measured Rydberg lifetimes and coherence times in tables 4.1 and 4.2 are plotted in figure 4.6. For the comparison with the measured lifetimes and coherence times, the $1/e$ times of the overlap functions η_i are used.

As can be seen in this figure, the Rydberg lifetimes increase with increasing n . The coherence times, on the other hand, do not seem to follow this behavior and, instead of increasing steadily, seem to first increase and then plateau. For the highest Rydberg state at $n = 96$, the coherence time even decreases. Factors contributing to the lower coherence time at $n = 96$ are the Rydberg–Rydberg interaction (see section 4.2.2.1), as well as outside electric fields destroying the stored excitation. The calculated longitudinal coherence time is slightly ν -dependent through Δk , giving rise to its curved trend, whereas




Figure 4.6: Calculated longitudinal dephasing time (τ_{long} , solid green) and transverse mode-overlap time (τ_{trans} , dashed magenta) as a function of the effective principal quantum number ν . The measured Rydberg-state lifetimes τ_{Ryd} (purple circles) are shown together with a ν^3 -scaling fit model (dashed blue-purple). The solid purple line shows τ_{tot} , the $1/e$ time of the product of all three overlap functions. The measured coherence times τ_c (cyan circles) are shown for comparison.

the calculated transverse mode-overlap time has no such ν -dependence and appears as a straight line. It is evident that both thermal mechanisms yield coherence-time limits far larger than what is experimentally observed, so they cannot be the dominant decoherence mechanism.

The measured coherence times reported here remain far below the calculated thermal limits of several hundred microseconds, confirming that thermal motion does not constrain the coherence times achieved in this work – a marked improvement over the thermal limit typical for Rb-87 [10, 39]. However, this improvement falls well short of what the thermal-limit theory alone would predict for Yb-174, indicating that an additional, as yet unidentified, decoherence mechanism must instead be dominant. This dominant dephasing mechanism is identified in the next section.

4.2.2.4 Collision-Induced Dephasing

As can be seen in the section above on thermal dephasing, this dephasing mechanism is not able to correctly explain the shortening of the coherence time compared to the measured Rydberg lifetime. In these models, one important factor has been explicitly neglected: the collision of atoms, primarily between a ground state atom and a Rydberg atom, which is a well-established dephasing channel in dense, ultracold Rydberg ensembles [47, 48]. Thus, a new model has to be introduced. Due to the complexity of this collisional interaction, in this thesis, a somewhat heuristic and simplified model, the hard-sphere model, is used to describe the collision-induced dephasing.

In this model, colliding atoms are treated as hard spheres, ignoring their internal structure and the interatomic potential. The hard-sphere model can be derived when considering the number of collisions an atom has in a given propagation distance. For such a particle in a homogeneous gas, and assuming that all particles travel linearly through space, that all particles behave as hard spheres, and that only two particles are involved in any given collision, the collision frequency can be described as [49]

$$\gamma_{\text{coll}} = \frac{(\text{Volume of Collisional Cylinder})(\text{Density})}{\text{Time}} = \sqrt{\frac{16\pi k_B T d^4 n_a^2}{m}}. \quad (4.25)$$

Here, the volume of the collisional cylinder – i.e., the cylinder that another atom has to enter for a collision to happen – is given by $V_{\text{coll}} = \pi d^2 \bar{v}_r t$, the cross-sectional area times the distance $\bar{v}_r t$ traveled at the mean relative speed, where d is the collision diameter, i.e., the sum of the two colliding atoms' radii. Since the ground-state atom's radius is negligible compared to the Rydberg orbital radius $R \propto v^2 a_0$, this reduces to the assumption $d \approx R$, so that the cross section is given by $\sigma = \pi R^2$, with a_0 the Bohr radius. Here, n_a is the atomic density, and m is the atomic mass of the identical collision partners. This collision rate yields a scaling for the dephasing rate with temperature and density as $\gamma_{\text{coll}} \propto n_a \sqrt{T}$.

This functional form matches the density- and velocity-averaged collision rate expression given in [49], and has been empirically confirmed to retain some validity in ultracold, dense Rydberg ensembles [47]. Nonetheless, it must be applied with caution, as at the present temperatures and densities, a regime is entered where many ground state atoms enter the Rydberg orbital volume simultaneously instead of colliding one at a time [48]. Still, in the following, the hard-sphere model is used to understand the temperature and density scaling of the coherence time.

To that end, the coherence time is measured for varying temperatures and densities (see section 2.1). Temperatures in the upper part of this range are accessed by varying the dipole trap laser power, which changes the trap potential depth and thus the number of atoms captured. In this trap, the temperature can be varied between 10 μK and 40 μK and the densities between $5 \times 10^{12} \text{ cm}^{-3}$ and $35 \times 10^{12} \text{ cm}^{-3}$. Even lower temperatures and higher densities can be reached by evaporating and recompressing the dipole trap. This means that the dipole trap power is adiabatically reduced to a lower intensity, such that the hotter atoms leave the trap. The trap is then recompressed by increasing the dipole trap power back to its original intensity, trapping the remaining colder atoms as a denser, colder cloud. With this scheme, typical temperatures of 3 to 10 μK and atomic densities of 5 to $60 \times 10^{12} \text{ cm}^{-3}$ can be reached. The atom number and thus the atomic density in both traps can be varied by changing the atom number in the green MOT loading into the dipole trap, enabling the large density scan in the dipole trap.

Using this ability of the experimental setup to reach these varying atomic conditions, a scan of temperature and density can be performed by recording the coherence time. This is done for the four different Rydberg states previously mentioned, $n = 42, 52, 73, 96$. The measured coherence rates are plotted versus the atomic conditions in figure 4.7 and are described by a model of the form

$$\gamma = \gamma_{\text{lt}} + \beta_{\text{coll}} n_a \sqrt{T}. \quad (4.26)$$

Here, γ_{lt} is the dephasing rate at zero temperature and density, i.e., the rate in the absence of any collisions, while β_{coll} has units of $\text{cm}^3 \text{ s}^{-1} \text{ K}^{-1/2}$ and describes all the factors in equation 4.25 aside from the density and temperature.

As can be seen in figure 4.7, the model reproduces the overall temperature and density dependence of the collision-induced dephasing over the range of measured atomic conditions reasonably well, yielding a $\chi_{\text{red}}^2 \approx 3.6\text{--}7.2$ across all Rydberg states. The coherence times as a function of the atomic conditions for the other Rydberg states are shown in the appendix A.

Using the results from figure 4.7, as well as the same model applied to the other Rydberg states, the collision-dephasing-independent rate γ_{lt} (at least within this model) can be determined. Since this




Figure 4.7: Measured coherence rates for $n = 52$ (circles) overlaid with the two-dimensional model shown in equation 4.26. The measured coherence rates and the model share the same colorbar, while the black lines denote its contour lines.

extracted rate is by construction independent of temperature and density, it should ideally reproduce the Rydberg decay rate. Thus, both the γ_{lt} rate and the Rydberg decay rate are plotted in figure 4.8.

In this figure, an estimate of the expected scaling can be made. Without the influence of the BBR, the Rydberg decay rate scales as $\gamma_0 \propto \nu^{-3}$ [15], due to the spontaneous emission rate introduced in section 4.1. The inverse-cubic scaling was determined for alkali atoms, but is nonetheless assumed to also hold for the effective principal quantum number of Yb-174 here. Rather than extracting the scaling exponent freely, the exponent is fixed to -3 for this comparison. Due to uncertainties in determining the Rydberg lifetime τ_{Ryd} , as well as γ_{lt} , the zero-temperature, zero-density limit of the coherence rate, neither data set is expected to reproduce the ν^{-3} scaling precisely. These imprecisions arise because the Rydberg lifetime measurement cannot distinguish which Rydberg state was ionized, and because the hard-sphere model used to extract γ_{lt} is itself only an approximate description of the collisional dephasing. Since the goal here is only a rough estimate of the expected scaling behavior, rather than a precise determination of it, this approximation nonetheless fulfills its purpose.

As a rough comparison, to give an image of how both the Rydberg decay rate and γ_{lt} , extracted from the collisional model, coincide with this scaling, the data for both are plotted together with a model function of the form $f(\nu) = a\nu^{-3}$; the model function for γ_{lt} further includes a constant offset. The additional offset captures ν -independent contributions not described by the ν -scaling itself. The resulting model parameters, shown in figure 4.8, show that the proportionality constant a for both the Rydberg decay rate and the extracted γ_{lt} agree within the error margins. The constant offset between the two determined rates can be attributed to further dephasing factors not captured by the collisional dephasing, or even additional factors in the experiment leading to an additional Rydberg decay or dephasing. This difference has not been quantified yet and remains to be precisely determined.

Aside from comparing the extracted coherence rate without collisions, the model can also be tested





Figure 4.8: Collision-dephasing-independent rate γ_{lt} (green circles) and Rydberg decay rate γ_{Ryd} (purple circles) as functions of the effective principal quantum number ν , together with the ν^{-3} model described in the text (solid lines in the respective colors).

regarding the extracted β_{coll} . This is done in figure 4.9.

The β_{coll} should scale with the cross section of the atoms as $\sigma = \pi R^2$, where R is the radius of the hard sphere. As known from [49], the electron orbital radius scales as $R \propto \nu^2 a_0$, where a_0 is the Bohr radius. Since the cross section scales as the radius squared ($\sigma \propto R^2$) and $R \propto \nu^2$, the expected scaling of β_{coll} is ν^4 . Describing the data with the power-law model function $f(\nu) = a\nu^p$ yields a scaling exponent of 3.351 ± 0.039 , which is slightly lower than expected. This can again be explained by the atomic cloud not fully satisfying the assumptions required for the hard-sphere model to be valid, an effect that was also observed in [48].

As was mentioned before, the cross section can be determined from β_{coll} . This extracted cross section is compared to the geometric cross section $\sigma_{geo} = \pi \langle r \rangle^2$, based on the mean Rydberg orbital radius, as well as to the ground-state–ground-state collisional cross section $\sigma_{gg} = 8\pi a^2$ [50], where $a = (5.55 \pm 0.08) \text{ nm} \approx 104.9 a_0$ is the s -wave scattering length for two ground-state Yb-174 atoms, as measured in [51]. This expression is the low-energy limit for identical bosons, valid when only s -wave scattering contributes, which is a good approximation at the μK temperatures probed here [50]. Together, these three cross sections are plotted in figure 4.9(b). It can be seen that the ground-state cross section is several orders of magnitude smaller than both σ_{eff} and σ_{geo} across the entire range of measured states, confirming the earlier assumption that ground-state–ground-state collisions can be neglected compared to collisions involving the Rydberg electron’s much larger orbit. The extracted σ_{eff} broadly tracks σ_{geo} , though not perfectly, – consistent with the extracted scaling exponent being slightly shallower than the expected ν^4 , as noted above.

In summary, it was found that the collisional dephasing is the dominant dephasing factor in the decoherence of the collective stored excitation. In particular, collisions between a ground state and a Rydberg atom are the dominant contributor, while ground-state–ground-state collisions, with a cross section σ_{gg} several orders of magnitude smaller than the extracted σ_{eff} , play no significant role. Despite



(a) Scaling of β_{coll} (circles) with the effective principal quantum number ν , expected to be ν^4 from the known scaling of the Rydberg collision cross section. The solid line shows the power-law model $\beta_{\text{coll}} \propto \nu^p$.

(b) Extracted effective cross section σ_{eff} (circles), theoretical geometric cross section $\sigma_{\text{geo}} = \pi \langle r \rangle^2$ (solid line), and ground-state-ground-state collisional cross section $\sigma_{\text{gg}} = 8\pi a^2$ (dashed line). Neither theory curve is a fit to the data.

Figure 4.9: Extracted scaling of β_{coll} (left) and the corresponding effective cross section, compared with theoretical predictions (right), both as functions of the effective principal quantum number ν .

the hard-sphere collisional model reproducing the temperature and density scaling of the dephasing rate well for all measured Rydberg states, this model is only a heuristic approach to describe the observed physics. In particular, at the densities and temperatures probed here, its assumption of well-separated, sequential collisions between only two atoms is not necessarily satisfied. This can be seen in the zero-temperature and zero-density lifetime extracted with this model, which shows a constant, ν -independent offset from the Rydberg lifetime measured in section 4.1.1. This offset has not been characterized yet, and the question remains whether it is caused by an additional dephasing mechanism, such as residual thermal dephasing, or an additional Rydberg decay in the experiment. Further, with this model, the cross section of the colliding Rydberg atom was determined and compared to the geometric cross section of a sphere with the mean Rydberg radius. While a similar scaling was found, with the extracted exponent somewhat below the expected ν^4 , an offset between both could be observed, likely because the atoms do not behave as hard spheres, in particular at these low temperatures.

4.2.3 Contribution of Rydberg Dimers

Besides the atomic Rydberg state, a Rydberg excitation in a dense atomic cloud can also form a bound molecular state with a nearby ground-state atom. For Yb-174, the first comprehensive study of such molecules was carried out only recently at this experiment [52]. For Rydberg-based storage of light, these molecules are relevant because, if a molecular state is excited alongside the atomic Rydberg state, the stored excitation evolves in two components with different energies, which directly modulates the retrieved signal. Understanding this contribution is therefore necessary to correctly interpret the measured storage and retrieval efficiency, and at the same time offers a way to access the molecular binding energies through the storage measurement itself.

Aside from the dephasing, one further effect was observed: a periodic revival of the storage and retrieval




Figure 4.10: Storage and retrieval efficiency for $n = 73$ at about $4 \mu\text{K}$ and a peak atomic density of $1 \times 10^{13} \text{ cm}^{-3}$, showing pronounced dimer oscillations. The measured data (green circles) are shown together with the model of equation 4.27 (black line).

efficiency after certain storage times. This effect is caused by the fact that, in addition to Rydberg atoms, such dimers can also be excited. As a result, both Rydberg atoms and dimers can exist at the same time in the atomic cloud [41, 52]. In Yb-174, this effect is less well documented and more complex to measure than in alkali atoms, due to the second valence electron and the more complex excitation spectrum.

To understand the scaling and to quantitatively reproduce the measured storage and retrieval efficiency curve, the two-species model introduced in section 4.2.1 can be used with an initial state of $|\psi(0)\rangle = \alpha|R\rangle + \beta|D\rangle$, with $|\alpha|^2 + |\beta|^2 = 1$, where the collective Rydberg excitation $|R\rangle$ and the collective dimer excitation $|D\rangle$ are defined in section 4.2.1. A fraction $n_R = |\alpha|^2$ of the stored excitation is thus stored as a collective Rydberg excitation, and a fraction $n_D = |\beta|^2$ as a collective dimer excitation. Both components evolve with their own energy and decay rate, and their energy difference is given by the dimer binding energy, $E_R - E_D = \hbar\Delta_D$ [41]. Upon retrieval, both contributions are converted into light in the probe mode and interfere. With the overall efficiency η_0 , the storage and retrieval efficiency is given by

$$\eta(t_s) = \eta_0 \left(n_R^2 \exp(-\gamma_{\text{Ryd}} t_s) + n_D^2 \exp(-\gamma_d t_s) + 2n_R n_D \exp\left(-\frac{\gamma_{\text{Ryd}} + \gamma_d}{2} t_s\right) \cos(2\pi\Delta_D t_s) \right), \quad (4.27)$$

where Δ_D is the dimer binding energy in frequency units, so that the retrieved signal oscillates at the frequency Δ_D .

The storage and retrieval efficiency measurements were performed in the same manner as in the previous sections.

Measuring these dimer oscillations in Yb-174 is a challenge because, for higher principal quantum numbers, the molecular resonances move closer to the atomic resonance and can start to overlap [52]. Over the range of the measured Rydberg states, precisely this effect can happen, and because many dimer and trimer peaks are present, more than one dimer resonance can contribute to the signal at the same time. In addition to multiple dimer contributions, higher-order effects cannot be neglected in this

measurement either. This makes extracting the individual dimer binding energies challenging.

Since the focus of this thesis lies elsewhere, this effect has not been measured in detail, but rather appeared as a byproduct of measuring the storage and retrieval efficiency at varying atomic conditions. Clearly quantifying this effect would therefore require further measurements, as well as some thought on how to improve the measurement apparatus to extract the different dimer and higher-order contributions. Nonetheless, a short overview of how such dimer oscillations can appear in the measured storage and retrieval efficiency can be given for the example state $n = 73$. The measured storage and retrieval efficiency curve for this state at roughly $4\text{ }\mu\text{K}$ and a peak atomic density of $1 \times 10^{13}\text{ cm}^{-3}$ is shown in figure 4.10.

As can be seen in this figure, the dimer oscillations are quite pronounced. This is normally not the case and cannot easily be reproduced at every Rydberg state. The measured dimer oscillation frequency Δ_D was found to be $(72.77 \pm 0.61)\text{ kHz}$. For comparison, an extrapolation of the ν -scaling of the binding energies of the outermost-well dimer Rydberg molecules yields $(107.66 \pm 0.04)\text{ kHz}$. The two values differ by a factor of about 1.5, which is reasonable given that several dimer resonances can contribute to the signal at the same time, and that the expected value is extrapolated from $n \leq 45$ to $n = 73$, so its uncertainty is larger than the given statistical uncertainty suggests, in particular if the ν^{-6} scaling does not hold exactly over this range [52].

In conclusion, the storage and retrieval efficiency measurements show a periodic revival that can be explained by the simultaneous excitation of Rydberg atoms and Rydberg dimers, whose coherent superposition beats at a frequency set by the dimer binding energy. For the example state $n = 73$, the extracted oscillation frequency of $(72.77 \pm 0.61)\text{ kHz}$ lies within a factor of about 1.5 of the value expected from the ν -scaling of the outermost-well dimers. This shows that dimer formation must be taken into account when interpreting storage measurements in dense Yb-174 clouds, as it can mask or mimic changes in the coherence time. At the same time, it demonstrates that the storage and retrieval efficiency itself can serve as a probe of Rydberg molecule binding energies, if the method is refined further. A systematic study across several Rydberg states, together with an improved setup that can resolve the individual dimer and higher-order contributions, would be required to turn this observation into a quantitative measurement.

Two-Photon Correlations

Having established in chapter 4 how long a collective Rydberg excitation (a stored polariton) retains its phase, the interaction between Rydberg polaritons propagating through the medium is examined in this chapter. The key property of Rydberg polaritons is that they map the strong interaction between Rydberg atoms onto photons. Via the Rydberg blockade, a polariton traveling through the atomic medium prevents a second photon from propagating under EIT conditions within the blockade radius, which renders the medium nonlinear on the level of single photons. Such photon–photon interactions are the basis for applications such as single-photon sources and photon–photon gates. Whether the blockade is effective in the given geometry, and thus whether this system is viable for such applications, is tested here directly. To this end, the photon statistics of the transmitted probe light are examined, quantified by the second-order correlation function $g^{(2)}$. It measures how likely the detection of a second photon is, given that a first photon has been detected. Since a single Rydberg excitation blockades the propagation of further photons, the detection of one photon is expected to suppress the detection of a second one shortly after. This effect, referred to as antibunching, cannot occur for any classical light field and thus directly reveals the nonlinearity on the level of single photons.

In this chapter, the theoretical background of the Rydberg blockade and the resulting antibunching is introduced in section 5.1, and the two-photon correlation function is defined in section 5.2. Building on this, the measured two-photon correlation function is presented in section 5.3. Under certain measurement conditions, a classical effect is observed in addition to the antibunching, which is attributed to ions created by the strong control light. This effect is characterized in section 5.4 using a minimal simulation, in which the mean detected photon rate changes once an ion is created in the medium.

5.1 Rydberg Blockade and Photon Antibunching

Under EIT conditions, introduced in section 4.2.1, the antibunching of the transmitted probe light is achieved via the Rydberg blockade. Each probe photon entering the atomic cloud is converted into a Rydberg polariton. The van der Waals interaction of the Rydberg component of this polariton shifts the Rydberg level of all atoms within the Rydberg blockade radius r_b out of resonance. The blockade radius is defined as the distance at which the van der Waals shift equals the width of the EIT transmission window, $2\hbar\Gamma_{\text{EIT}}$ (see equation 2.2). Since $C_6 \propto \nu^{11}$, the blockade radius scales as $r_b \propto \nu^{11/6}$ at constant Γ_{EIT} [15].

Within the blockade radius, the [EIT](#) condition for a second polariton is broken and the medium acts as a resonant two-level medium. A second probe photon entering the blockade radius of such a collective Rydberg excitation is then absorbed on the $|g\rangle \rightarrow |e\rangle$ transition and re-emitted by spontaneous emission into a random spatial direction, meaning that it is not collected into the detected mode.

In this work, full blockade is aimed to be achieved by choosing an appropriate Rydberg state at which the Rydberg blockade radius exceeds the probe beam waist, $r_b > w_0$. Under these conditions, the entire transversal probe beam can be blocked by a single Rydberg excitation, and no second polariton can pass this preexisting polariton. Two photons are consequently only transmitted if they are separated in time by at least the time the first polariton requires to propagate through the blockade radius with the group velocity v_g of the slow-light polariton, $\tau_b \approx r_b/v_g$ (see equation [4.7](#)).

5.2 Second-Order Correlation Function

The photon statistics of the probe photons transmitted through the atomic cloud as Rydberg polaritons, i.e., not scattered out of the probe beam via spontaneous emission, are quantified by the two-photon correlation function $g^{(2)}(t_1, t_2)$. This section is based on [\[53\]](#). The two-photon correlation function is based on a [HBT](#) measurement (see section [2.4](#)), in which the transmitted probe light is split by a beam splitter, and each output is directed onto a separate detector. Two detectors are required, since the dead time of a single detector, on the order of tens of nanoseconds, prevents it from registering two photons in close succession. Intuitively, $g^{(2)}(t_1, t_2)$ gives the probability of detecting a second photon at t_2 on one detector, given that a photon was already detected at t_1 on the other, normalized to the photon detection events on both detectors at the respective times, taken independently. The four [SPCMs](#) in the setup are used for averaging, i.e., the two-photon correlation function is calculated between every pair of the four detectors and averaged. This improves the statistics, since the probability that two photons are registered on different detectors increases from 1/2 for two detectors to 3/4 for four detectors.

Depending on whether the detection of a first photon increases, decreases, or does not affect the probability of detecting a second one, three cases can be distinguished. For $g^{(2)} = 1$, the detection events are uncorrelated, which corresponds to a coherent light source such as a laser. Detecting one photon provides no information about detecting the next one. For $g^{(2)} > 1$, generally referred to as bunching, photons are detected close to each other more often than randomly. When one photon is detected, the chance of detecting another close in time is thus increased. This is the case, e.g., for thermal light. For $g^{(2)} < 1$, referred to as antibunching, when one photon is detected, the probability of detecting a second is decreased. This is a completely nonclassical effect: since no classical intensity can fluctuate less than the constant intensity of coherent light, $g^{(2)}(0) < 1$ cannot be explained by any classical light field and is a direct signature of the nonlinearity imprinted on the transmitted light. The Rydberg blockade produces exactly this effect: in the limit of a perfectly blocked probe beam, at most one photon is transmitted within the time corresponding to the Rydberg blockade radius. The width of the antibunching dip thus reflects the blockade radius, while its depth reflects how completely the probe beam is blocked [\[17\]](#).

Classically, the correlation between the intensities I_1 and I_2 recorded at the two detectors can be characterized via

$$g^{(2)}(t_1, t_2) = \frac{\langle I_1(t_1)I_2(t_2) \rangle}{\langle I_1(t_1) \rangle \langle I_2(t_2) \rangle}, \quad (5.1)$$

where $\langle \cdot \rangle$ denotes the average over many realizations, here over many probe pulses. For equal times, both outputs are proportional to the same input intensity $I(t)$, such that $g^{(2)}(t, t) = \langle I(t)^2 \rangle / \langle I(t) \rangle^2 \geq 1$ holds for any classical field, since the variance of the intensity cannot be negative.

In the quantum picture of light, the intensities are replaced by the photon number operators $\hat{n}_i$ of the two detector modes. The classical definition then becomes

$$g^{(2)}(t_1, t_2) = \frac{\langle \hat{n}_1(t_1) \hat{n}_2(t_2) \rangle}{\langle \hat{n}_1(t_1) \rangle \langle \hat{n}_2(t_2) \rangle}. \quad (5.2)$$

Experimentally, the expectation values are replaced by averages over all probe pulses of the photon counts $n_i(t)$ recorded by detector i in time bins of 50 ns. For every pair of distinct detectors $i < j$, the correlation function


$$g_{ij}^{(2)}(t_1, t_2) = \frac{\langle n_i(t_1) n_j(t_2) \rangle}{\langle n_i(t_1) \rangle \langle n_j(t_2) \rangle} \quad (5.3)$$

is calculated. The resulting six correlation functions are averaged and symmetrized with respect to $t_1 \leftrightarrow t_2$. The one-dimensional correlation function $g^{(2)}(\tau)$ is obtained by averaging the numerator and the denominator of equation 5.3 separately along lines of constant $\tau = t_1 - t_2$ before dividing.

5.3 Experimentally Determined Second-Order Correlation Function

With the Rydberg blockade and the two-photon correlation function established in sections 5.1 and 5.2, it remains to be tested whether an effective blockade of the probe beam is possible for this setup. To this end, the two-photon correlation function of the transmitted probe light is determined first under conditions of full blockade. Subsequently, measurements over longer probe pulses and for varying atomic conditions are presented, in which an additional structure in the correlation function becomes visible.

In order to achieve full blockade as described in section 5.1, the $n = 97$ Rydberg state is chosen. The two-photon correlation function is determined according to section 5.2. This is done both with atoms present, in the signal pulse, and without atoms, in a reference pulse, to obtain a comparable baseline. The resulting two-photon correlation functions are shown in figure 5.1.



(a) Measured two-photon correlation function $g^{(2)}(t_1, t_2)$ at $n = 97$, showing the clean antibunching dip along the $\tau = 0$ diagonal. The photon rates are plotted in the side panels.

(b) Corresponding reference measurement without atoms present, showing a correlation function that is constant at one. The photon rates are plotted in the side panels.

Figure 5.1: Two-photon correlation function at $n = 97$ with and without atoms in figures 5.1(a) and 5.1(b), respectively, measured with $N_a = 2.6 \times 10^5$ atoms at $T = 12.3 \mu\text{K}$ and a control power of $P_c = 70 \text{ mW}$, i.e., under the same conditions as the EIT window and slow-light delay in figure 4.4.




Figure 5.2: One-dimensional two-photon correlation function $g^{(2)}(\tau)$, dependent on the difference $\tau = t_1 - t_2$, measured at $n = 97$, which is obtained from the two-dimensional plot in figure 5.1(a) via averaging along lines parallel to the $\tau = 0$ diagonal. The signal (with atoms) and reference (no atoms) correlation functions are plotted in dark and light green, respectively, and the model of equation 5.4 describing the signal as a black line.

At $n = 97$, a pronounced antibunching dip along the $\tau = 0$ diagonal can be seen in figure 5.1(a). Away from this feature, for time differences well above the width of the dip, the two-photon correlation function remains close to one. The photon rate, shown in the side panels, remains relatively constant along the


pulse duration, with only a slight decline towards later times, apart from the rise at the beginning, which follows the rising edge of the probe pulse. The time axes of the signal and reference measurement are identical. The signal pulse, however, arrives later than the reference pulse due to the slow-light delay in the atomic medium (see equation 4.7). The reference measurement without atoms in figure 5.1(b) is constant at one, as expected for coherent light.

The one-dimensional correlation function in figure 5.2 is obtained as described in section 5.2. In order to extract the depth and width of the antibunching dip, the phenomenological Gaussian function

$$g^{(2)}(\tau) = g_{\infty} - (g_{\infty} - g_0) \exp\left(-\ln 2 \frac{\tau^2}{t_{\text{HWHM}}^2}\right) \quad (5.4)$$

is used to describe the signal correlation function, which yields $g^{(2)}(0) = 0.145 \pm 0.018$, $t_{\text{HWHM}} = (0.611 \pm 0.015) \mu\text{s}$, and a baseline of $g_{\infty} = 1.038 \pm 0.004$. This measurement shows that the Rydberg blockade produces a clean, deep antibunching dip. The value of $g^{(2)}(0)$ means that the probability of detecting two photons within the same time bin is reduced to about 15 % of that expected for uncorrelated photons at the same mean rate, while the width of the dip reflects the time τ_b for which a transmitted polariton blockades the medium (see section 5.1). The residual value above zero can be attributed to an incomplete blockade of the probe beam and to uncorrelated background counts. For longer probe pulses at the $n = 73$ Rydberg state, however, additional structure appears in the correlation function.

An example of such a measurement, taken at $n = 73$ with a control power of roughly 60 mW, is shown in figure 5.3.



(a) Measured two-photon correlation function for photons passing through the atomic cloud at $n = 73$, a control power of roughly 60 mW, a temperature of around 23 μK , and a density of around $30 \times 10^{12} \text{ cm}^{-3}$. Along the $\tau = 0$ diagonal, the antibunching can be seen, and the panels alongside the two-photon correlation function show the mean photon rate trend qualitatively.

(b) Measured two-photon correlation function for photons in the reference pulse without the atomic cloud. There is no antibunching dip, and the correlation function is constant at 1, which is to be expected for coherent light. Again, the panels alongside the correlation function show the probe pulse shape.

Figure 5.3: Two-photon correlation function for the probe beam passing through the atomic cloud in figure 5.3(a) and in the reference pulse without atoms in figure 5.3(b).

In the correlation function with the atomic cloud present (figure 5.3(a)), the antibunching dip along the $\tau = 0$ diagonal is still visible, but it is surrounded by a bunching feature, $g^{(2)} > 1$, which becomes more pronounced towards the end of the pulse. At the beginning of the pulse, the correlation at $\tau = 0$ falls to roughly 0.8. Towards the end of the pulse, it rises above one, but remains lower than the bunching on both sides of the diagonal, such that the antibunching dip is still visible as a local minimum.

For comparison, the corresponding reference measurement is shown in figure 5.3(b), which is constant at one with a constant photon rate, in contrast to the signal pulse.

This bunching feature is attributed to the ionization of ground-state atoms by the strong control light, introduced in section 4.2.2.2. This attribution is supported by a measurement of the two-photon correlation function in which an electric field is switched on for a short interval during the probe pulse, which removes the created ions from the atomic cloud but also temporarily shifts the medium out of the EIT condition. After the electric field is switched off again, both the photon rate and the antibunching dip are recovered, which indicates that both effects are caused by charged particles in the medium rather than by the electric field itself.

In order to investigate how this feature depends on the atomic medium, the measurement was repeated for different atomic conditions. The effect is present for all investigated atomic conditions, at densities from roughly 20×10^{12} to $40 \times 10^{12} \text{ cm}^{-3}$ and temperatures from 12 to 31 μK . In order to model its shape, a minimal simulation is introduced in section 5.4.

In summary, the measured two-photon correlation function demonstrates that the Rydberg blockade is effective in the given geometry. At $n = 97$, where the blockade radius is expected to exceed the probe beam waist, the transmitted light is strongly antibunched with $g^{(2)}(0) = 0.145 \pm 0.018$ and a half width at half maximum of $t_{\text{HWHM}} = (0.611 \pm 0.015) \mu\text{s}$, confirming that a single Rydberg excitation largely blockades the probe beam. For longer probe pulses at $n = 73$, however, a bunching feature develops around the antibunching dip, which grows over the course of the pulse and eventually raises the correlation at the dip above one. This feature is observed for all investigated atomic conditions and is attributed to the ionization of ground-state atoms by the control light, which is examined in more detail in the following section.

5.4 Ion Simulation

As shown in section 5.3, the measured two-photon correlation function exhibits a bunching feature around the antibunching dip, which is attributed to ions created by the control light. In order to test this attribution, a minimal model is introduced in this section, in which photon counts are simulated under two assumptions: an ion is created at a random time, and its creation changes the photon rate abruptly. The effect of a second ion within the same pulse and the Rydberg blockade itself are not included in the model. If this model, based on the experimentally measured ion and photon count rates, reproduces the bunching feature observed in figure 5.3(a), this gives a strong indication of the ions being the origin of the bunching feature.

As established in section 4.2.2.2, ions are created when a ground-state atom in the medium is ionized by the absorption of two control field photons, which is energetically possible since the energy of two control photons exceeds the ionization energy of Yb-174. The model used to simulate the feature around the antibunching dip is based on the ion being created at a random time, drawn from a Poisson process

with constant rate. Only the first ion of each pulse is taken into account, and it decreases the detected photon rate exactly at the time of its creation. Since this switching time affects all following detected photons, regardless of which detector registers them, it correlates counts recorded at two different times, even though the single-photon detection events are independent of one another. This leads to a visible bunching in the correlation function.

The ion is created at a random time in each pulse. Therefore, the time at which the photon rate drops varies from pulse to pulse. At a given time, the ion has already been created in some pulses, such that the photon rate is low, while in the other pulses, no ion has been created yet, such that the photon rate is still high. A detected photon therefore most likely stems from a pulse in which no ion has been created yet, since more photons are detected in these pulses.

If a second photon is detected shortly after the first one, it is unlikely that the ion was created in between. The photon rate is then still high, such that the second photon is detected with a higher probability than on average over all pulses. This results in $g^{(2)} > 1$ close to the diagonal $t_1 = t_2$. If the two photons are detected further apart, it becomes more likely that the ion was created in between, which abruptly reduces the photon rate for the second photon. The probability of detecting the second photon then approaches the average over all pulses, and the correlation decreases with increasing time difference.

Moreover, the bunching increases over the course of the pulse. At early times, an ion has been created in almost none of the pulses, such that the photon rate is the same in all pulses and the detection of a photon carries no information about the pulse. Since the probability that the ion has already been created increases over time, the pulses differ more and more in their photon rate at later times. Detecting a photon then selects pulses with a high photon rate more strongly, which increases the bunching.

While the detected photon rate before the creation of the ion matches the experimentally measured photon rate, after the ion has been created, the photon rate is reduced to a fraction $f \in [0, 1]$ of its initial value. This reduction of the photon rate reflects that an ion in the medium shifts the Rydberg level of nearby atoms out of the two-photon resonance via the Stark shift of its electric field, breaking the EIT condition in this region, and thus the transmitted photon rate is suppressed. The underlying rates are the ionization rate μ_{ion} and the pre-jump photon rate μ_{photon} . In the simulation, the following steps are carried out for each pulse:

1. for each time bin of 50 ns, the number of created ions is drawn from a Poisson distribution with mean $\mu_{\text{ion}} \Delta t$, with the first bin containing an ion defining the time of the rate change,
2. the photon rate is set to μ_{photon} before and to $f\mu_{\text{photon}}$ after this time, and
3. for each of the four detectors, the photon counts in each time bin are drawn from a Poisson distribution with this rate.


These steps are repeated for 10^5 pulses, and the two-photon correlation function is calculated from the simulated counts in the same way as from the measured ones (see equation 5.3).

The ionization rate is obtained from the number of ions detected per pulse, corrected for the overall ion detection efficiency η , and is taken to be constant over the pulse. For the energy of the arriving ions of slightly above 2 keV, the detection efficiency of the MCP¹ is specified to be between 60 % and 85 % [25]. Further factors reduce the detection efficiency, such as its dependence on the angle of incidence of the

¹ Hamamatsu F4655-13

ions and its degradation over the time of use. Since the actual detection efficiency is not known, it was estimated as $\eta = 0.4$ for the simulation. All simulation parameters are listed in table A.1 in the appendix. The pre-jump photon rate is chosen such that the simulated mean photon rate matches the measured one over the duration of the probe pulse. Using these rates of creating an ion and of detecting a photon before and after the ion event, detector clicks can be simulated resembling the measured data. These simulated clicks are then compared against the measured correlation function, using both the two-dimensional $g^{(2)}(t_1, t_2)$ map and the mean photon rate over the pulse, shown in the side panels. The photon factor is chosen such that the mean squared difference between the simulated and measured two-dimensional correlation maps is minimal, within the window from 5 to 45 μs , excluding the $\tau = 0$ diagonal.

The photon data are simulated for three data sets with different atomic conditions at the same control power of roughly 60 mW at the $n = 73$ Rydberg state. Since the number of created ions depends on the atomic conditions, the strength of this feature varies between the data sets. For the data set taken at an atomic density of $30 \times 10^{12} \text{ cm}^{-3}$ and a temperature of 23 μK , the two-photon correlation function obtained from the simulated data is shown in figure 5.4 together with the difference between the simulated and measured $g^{(2)}(t_1, t_2)$ function.



(a) Two-photon correlation calculated from the simulated photon counts with an ionization event changing the mean photon rate at a random time at most once in each pulse. In the panels alongside the correlation plot, the mean photon rates are shown for the correlated times.

(b) Difference between the simulated and measured two-photon correlation function, $g_{\text{sim}}^{(2)} - g_{\text{meas}}^{(2)}$. Since the simulation does not have an antibunching dip, the diagonal appears as a positive difference. Closer to the diagonal, the correlation is underestimated and thus negative, while further away, it is overestimated and thus positive.

Figure 5.4: Two-photon correlation function obtained from the simulation in figure 5.4(a), as well as the difference between the correlation functions calculated from the simulated and measured (see figure 5.3(a)) detected photons in figure 5.4(b). The data is simulated for the data set with 0.153 detected ions per pulse, corresponding to an ionization rate of $\mu_{\text{ion}} = 0.0106 \mu\text{s}^{-1}$ for $\eta = 0.4$, a pre-jump photon rate of $\mu_{\text{photon}} = 0.356 \mu\text{s}^{-1}$ per detector, and $f = 0.11$. The parameters of all data sets are listed in table A.1.

As can be seen in this figure, the simulated two-photon correlation function only approximately recreates the measured two-photon correlation function. There is no antibunching dip, since the blockade is not included in the model. Aside from that, the simulated correlation function decreases more slowly than the measured correlation moving away from the diagonal. Further, the measured correlation also peaks

higher closer to the antibunching dip. This is visualized in figure 5.4(b), where the difference between the simulated $g_{\text{sim}}^{(2)}$ and measured $g_{\text{meas}}^{(2)}$ correlation functions is shown.

As was previously mentioned, the correlation function was calculated for varying atomic conditions at the same control power. Higher or lower atomic density and temperature affect the OD of the medium. Since the ion breaks the EIT condition in its vicinity, the probe light is absorbed there on the two-level transition, which is expected to reduce the transmitted photon rate more strongly at higher OD. For the three data sets, photon factors of $f = 0.18, 0.11$, and 0.06 are found, in the order of increasing atomic density and temperature (see table A.1). In all three cases, the photon rate is thus strongly decreased after the creation of an ion, and the decrease becomes stronger with increasing atomic density and temperature. The photon factors vary with the assumed detection efficiency, as can be seen in table A.1.

Overall, the simulation reproduces the bunching around the $\tau = 0$ diagonal and its increase over the course of the pulse, i.e., the qualitative structure of the measured correlation function apart from the antibunching dip. The remaining quantitative differences can likely be attributed to the simplicity of the model. It does not consider the effect of two ions on the medium, nor does it include the antibunching dip. This supports the hypothesis that the structure in the two-photon correlation function can indeed be attributed to the ions created by the control field in the medium.

Conclusion

In this thesis, light propagation and Rydberg polariton dynamics are studied in a dense ultracold ensemble of [Yb-174](#) atoms. The three addressed aspects are

1. the [OD](#) of the atomic cloud,
2. the Rydberg lifetime and the coherence time of collective Rydberg excitations, and
3. photon–photon interactions mediated by the Rydberg blockade.

In chapter 2, to facilitate effective photon–photon interactions, the experimental apparatus is extended by motorized mirrors for dipole trap alignment, improving the positioning accuracy of the atomic cloud from approximately $15\text{ }\mu\text{m}$ to $2\text{ }\mu\text{m}$. Additionally, the probe beam focus position is located with two independent methods. Well-defined probe pulse shapes are realized via an [AOM](#) calibration, and the control beam is reshaped into a flat-top beam profile at the probe beam focus position. With these upgrades, the atomic cloud can be aligned reproducibly to the probe beam focus without manual access to the beam path, atom loss due to the control beam is reduced, and the precisely shaped probe pulses required for the coherence time measurements are available.

In chapter 3, the [OD](#) predicted by the conventional one-dimensional model, which is commonly used to extract the [OD](#) from the detuning-dependent probe transmission, is shown to overestimate the [OD](#) experienced by the probe beam in the dense clouds prepared in this apparatus. Calculated along the probe beam axis from the density distribution of the atomic cloud, this [OD](#) exceeds the [OD](#) simulated with the [ASM](#) by 26 % at a peak atomic density of $9.7 \times 10^{12}\text{ cm}^{-3}$ and by a factor of about 3 at $2.5 \times 10^{13}\text{ cm}^{-3}$, and the deviation grows further with increasing density. The reason is that the probe beam is diffracted and refracted by the transversally inhomogeneous cloud, so that part of the light bypasses the densest region of the cloud – an effect neglected in the one-dimensional model. Agreement between both approaches is only recovered in the limit of a cloud that is transversally large compared to the probe beam. The [OD](#) obtained directly from the transmitted photon rate follows the density scaling of the simulated [OD](#), but lies a factor of 2 to 3 below it, which is attributed to the non-Gaussian shape of the real probe beam. Hence, neither experimental extraction method yields the [OD](#) experienced by the probe beam. The transmitted probe beam shape, simulated with the [ASM](#) by propagating the probe field through the atomic density calculated from the dipole trap potential, qualitatively agrees with an independent camera-based measurement of the [RQO](#) group. Consequently, transmission measurements on dense

clouds cannot be converted into an OD without accounting for the beam reshaping by the inhomogeneous atomic cloud. For this apparatus, the simulated OD is to be taken as the more reliable value.

In chapter 4, the coherence time of a stored collective Rydberg excitation is found to lie between 10 μs and 17 μs , significantly below the Rydberg lifetime, and to be limited by collisions between Rydberg and ground-state atoms rather than by thermal motion. The measured Rydberg lifetimes, which also contain contributions of Rydberg states populated by BBR owing to the non-state-selective detection, exceed the values calculated with PairInteraction [22, 23] by 9 to 21 % for $n = 42, 52$, and 73, and by a factor of 2.4 for $n = 96$. Owing to the similar wavelengths of the two excitation lasers in Yb-174, thermal motion, i.e., longitudinal dephasing and the transversal mode overlap, only limits the coherence time to several hundred microseconds and is therefore not the dominant mechanism, in contrast to typical alkali-atom platforms. Instead, the temperature and density dependence of the dephasing rate is reproduced for all measured Rydberg states by a hard-sphere model of collisions between Rydberg and ground-state atoms, with the extracted cross section scaling as $v^{3.35}$, slightly below the geometrically expected v^4 . Long storage times are therefore only reached at low atomic densities and temperatures, which competes with the high densities required for a large OD. Additionally, a periodic revival of the storage and retrieval efficiency is observed at $n = 73$ and attributed to Rydberg dimers. This provides a first indication that dimer binding energies at high principal quantum numbers can be accessed through coherence time measurements.

In chapter 5, the Rydberg blockade is shown to mediate an effective photon–photon interaction in the given geometry. At $n = 97$, where the Rydberg blockade radius is expected to exceed the probe beam waist, the transmitted light is strongly antibunched with $g^{(2)}(0) = 0.145 \pm 0.018$ and a half width at half maximum of $t_{\text{HWHM}} = (0.611 \pm 0.015) \mu\text{s}$. For longer probe pulses at $n = 73$, a bunching feature develops around the antibunching dip and eventually lifts it above $g^{(2)} = 1$. This feature is attributed to ionization of ground state atoms by the control light, which is supported by the recovery of both the photon rate and the antibunching dip after the ions are removed by an electric field pulse applied during the probe pulse. A minimal model, in which an ion abruptly changes the transmitted photon rate at a random time, reproduces its overall shape qualitatively. This ionization-induced structure therefore limits two-photon correlation measurements for longer probe pulses.

The results of this thesis open up several directions for future research. A full characterization of the probe beam, including aberrations, or direct imaging of the transmitted probe beam after the atomic cloud would improve the accuracy of the determined OD. A collisional model beyond the hard-sphere approximation could clarify the remaining offset between the collision-free coherence rate and the Rydberg decay rate. A systematic study of the dimer revivals across several Rydberg states would provide access to dimer binding energies at higher-lying Rydberg states, where the molecular resonances overlap with the atomic resonance and can no longer be resolved spectroscopically. Further, the model describing the bunching caused by ionization through the control light could be extended to include the antibunching, so that the full measured correlation function could be reproduced quantitatively.

Abbreviations

AOM Acousto-Optic Modulator. [17–19](#), [71](#)

ASM Angular Spectrum Method. [2](#), [23](#), [24](#), [26–30](#), [33–36](#), [38–40](#), [71](#)

BBR Black Body Radiation. [42–45](#), [57](#), [72](#), [75](#), [76](#)

EIT Electromagnetically Induced Transparency. [1](#), [9](#), [10](#), [23](#), [46–49](#), [51](#), [62](#), [63](#), [65](#), [67](#), [68](#), [70](#)

HBT Hanbury Brown Twiss. [20](#), [21](#), [63](#)

MCP Microchannel Plate. [2](#), [5](#), [20](#), [21](#), [43](#), [68](#)

MOT Magneto-Optical Trap. [2–5](#), [9–12](#), [14–16](#), [56](#)

NQO Nonlinear Quantum Optics. [2](#)

OD Optical Depth. [2](#), [3](#), [6](#), [10–12](#), [14–16](#), [21–24](#), [30](#), [31](#), [33–41](#), [70–72](#)

PI Physik Instrumente. [6](#)

Rb-87 Rubidium-87. [1](#), [39](#), [51](#), [52](#), [55](#)

RF Radio Frequency. [17](#), [18](#)

RQO Rubidium Quantum Optics. [23](#), [39](#), [40](#), [71](#)

SPCM Single-Photon Counting Module. [5](#), [18](#), [21](#), [37–39](#), [46](#), [63](#)

Yb-174 Ytterbium-174. [1–3](#), [9](#), [23](#), [51](#), [52](#), [55](#), [57–61](#), [67](#), [71](#), [72](#)

YQO Ytterbium Quantum Optics. [3](#), [39](#), [40](#), [80](#)

Supplementary Material

A.1 Probe Beam Tracing

In addition to the vertical probe beam trace discussed in chapter 2, the probe beam trace was further recorded with the horizontal imaging. The measurement and evaluation were carried out analogously to the vertical trace. The horizontal probe beam trace is shown in figure A.1, and the beam waist extracted from it is shown in figure A.2.




Figure A.1: Horizontal probe beam trace and slices transversal to the probe beam, showing the beam profile at varying points along its propagation direction. The white dashed lines in the upper plot indicate the positions of the transversal slices shown in the lower plots. In the lower plots, the averaged signal (blue circles) is shown together with the double-Gaussian model of equation 2.3 (black line).






Figure A.2: Beam waist extracted from the horizontal probe beam trace. The extracted beam widths are shown as circles with error bars, and the Gaussian beam profile of equation 2.4 describing them as a black line. As for the vertical trace, the extracted waist is larger than the actual probe beam waist, due to the saturation of the atomic response at high probe intensities.

A.2 Black Body Radiation and Spontaneous Emission Rates


The transition rates from the Rydberg states with $n = 52, 73$, and 96 to other states, calculated with PairInteraction [22, 23] analogously to the $n = 42$ state in chapter 4, are shown in figures A.3(a) to A.3(c). For all three states, the single largest rate is given by spontaneous emission to a low-lying state, while BBR mainly couples to states with principal quantum numbers close to that of the initially excited state. The BBR transition rate increases relative to the spontaneous emission transition rate for increasing effective principal quantum numbers.



(a) $n = 52$.



(b) $n = 73$.



(c) $n = 96$.

Figure A.3: **BBR** and spontaneous emission transition rates to other states from the Rydberg states with $n = 52$, 73, and 96. The histogram columns show in green the transition rates generated by spontaneous emission and in purple the transition rates for **BBR** at room temperature 300 K.

A.3 Collisional Dephasing Maps for $n = 42, 73$, and 96

For completeness, the two-dimensional representations of the coherence decay rate γ_c as a function of the temperature T and the atomic density n_a are shown for the principal quantum numbers $n = 42, 73$, and 96 . They were obtained in the same way as the representation for $n = 52$ in figure 4.7.




Figure A.4: Coherence decay rate γ_c as a function of the temperature T and the atomic density n_a for (a) $n = 42$, (b) $n = 73$, and (c) $n = 96$. The color map shows the model of equation 4.26, evaluated with the parameters extracted by orthogonal distance regression, and the black lines are contours of constant γ_c . The circles mark the measured decay rates, colored on the same scale, so that deviations from the model appear as a color contrast to the background. The extracted parameters are given below each plot.

A.4 Ion Simulation

The ion simulation introduced in chapter 5 was carried out for three data sets with different atomic conditions at the same control power of roughly 60 mW at the $n = 73$ Rydberg state. The data set

at a density of $30 \times 10^{12} \text{ cm}^{-3}$ and a temperature of $23 \mu\text{K}$ is discussed in the main text. For the two remaining data sets, the measured two-photon correlation functions are shown in figure A.5, the correlation functions calculated from the generated data in figure A.6, and the difference between both in figure A.7. The atomic conditions and simulation parameters of all three data sets are listed in table A.1.




Figure A.5: Measured two-photon correlation function $g^{(2)}(t_1, t_2)$ for the two remaining data sets. The corresponding simulation parameters are listed in table A.1.





Figure A.6: Two-photon correlation function calculated from generated data for the two remaining data sets. The data is generated with the parameters listed in table A.1.


 (a) Density of $20 \times 10^{12} \text{ cm}^{-3}$ and temperature of 12 μK .

 (b) Density of $40 \times 10^{12} \text{ cm}^{-3}$ and temperature of 31 μK .

 Figure A.7: Difference $g_{\text{sim}}^{(2)} - g_{\text{meas}}^{(2)}$ between the correlation functions calculated from the simulated and measured detected photons for the two remaining data sets.

 Table A.1: Atomic conditions and parameters of the ion simulation for the three data sets, taken at the same control power of roughly 60 mW at $n = 73$. In each case, 10^5 pulses of 1 100 time bins of 50 ns (55 μs) were generated for four detectors, and the correlation functions were rebinned to 1 μs . The pre-jump photon rate μ_{photon} per detector was chosen such that the simulated mean photon rate matches the measured one on the plateau of the probe pulse (8 to 42 μs). The ionization rate μ_{ion} was calculated from the number of detected ions per pulse N_{ion} and the plateau length of the probe pulse, for an overall ion detection efficiency of $\eta = 0.4$. The photon factor f was adjusted to the measured two-dimensional correlation functions within 5 to 45 μs , excluding the diagonal; the mean squared error (MSE) of this comparison is given as well. The range of f in brackets corresponds to $\eta = 0.5$ and $\eta = 0.3$.

n_a (10^{12} cm^{-3})	T (μK)	μ_{photon} (μs^{-1})	N_{ion}	μ_{ion} (μs^{-1})	f	MSE (10^{-3})
20	12	0.529	0.0568	0.0040	0.18 (0.09–0.28)	0.46
30	23	0.356	0.1532	0.0106	0.11 (0.02–0.19)	2.9
40	31	0.336	0.1492	0.0103	0.06 (0–0.16)	3.6

Acknowledgements


First of all, I would like to thank Prof. Hofferberth for giving me the opportunity to write my thesis at the [YQO](#) experiment in the Nonlinear Quantum Optics group, and Prof. Wang for agreeing to be the second corrector of this thesis.

A big thank you goes to Tangi, who explained so much to me and never got tired of the many long discussions we had while working on all of this.

Thanks also to Xin and Eduardo for their help in the lab and for always taking the time to answer my many questions. Xin found every problem in the setup remarkably fast and made it possible to get things done just as quickly – especially changing the Rydberg state, where he is still the unbeaten champion. Finally, thanks to the whole group for the conversations and, of course, for the cake. Whether for a published paper, a birthday, a submitted thesis or simply because someone felt like it, there was always a reason to bring cake – and I will happily keep up this tradition when handing in.



The Nonlinear Quantum Optics group.



The Ytterbium Quantum Optics lab.

Bibliography

- [1] J. L. O’Brien, A. Furusawa and J. Vučković, *Photonic quantum technologies*, *Nature Photonics* **3** (2009) 687, ISSN: 1749-4893, URL: <http://dx.doi.org/10.1038/nphoton.2009.229> (cit. on p. 1).
- [2] N. Gisin and R. Thew, *Quantum communication*, *Nature Photonics* **1** (2007) 165, ISSN: 1749-4893, URL: <http://dx.doi.org/10.1038/nphoton.2007.22> (cit. on p. 1).
- [3] A. I. Lvovsky, B. C. Sanders and W. Tittel, *Optical quantum memory*, 2010, arXiv: [1002.4659](https://arxiv.org/abs/1002.4659) [quant-ph], URL: <https://arxiv.org/abs/1002.4659> (cit. on p. 1).
- [4] D. E. Chang, V. Vuletić and M. D. Lukin, *Quantum nonlinear optics – photon by photon*, *Nature Photonics* **8** (2014) 685 (cit. on p. 1).
- [5] B. Hacker, S. Welte, G. Rempe and S. Ritter, *A photon–photon quantum gate based on a single atom in an optical resonator*, *Nature* **536** (2016) 193, ISSN: 1476-4687, URL: <http://dx.doi.org/10.1038/nature18592> (cit. on p. 1).
- [6] D. Tiarks, S. Schmidt-Eberle, T. Stolz, G. Rempe and S. Dürr, *A photon–photon quantum gate based on Rydberg interactions*, *Nature Physics* **15** (2018) 124, ISSN: 1745-2481, URL: <http://dx.doi.org/10.1038/s41567-018-0313-7> (cit. on p. 1).
- [7] K. M. Birnbaum et al., *Photon blockade in an optical cavity with one trapped atom*, *Nature* **436** (2005) 87 (cit. on p. 1).
- [8] P. Lodahl, S. Mahmoodian and S. Stobbe, *Interfacing single photons and single quantum dots with photonic nanostructures*, *Reviews of Modern Physics* **87** (2015) 347, ISSN: 1539-0756, URL: <http://dx.doi.org/10.1103/RevModPhys.87.347> (cit. on p. 1).
- [9] O. Firstenberg, C. S. Adams and S. Hofferberth, *Nonlinear quantum optics mediated by Rydberg interactions*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **49** (2016) 152003, ISSN: 1361-6455, URL: <http://dx.doi.org/10.1088/0953-4075/49/15/152003> (cit. on pp. 1, 14, 23, 24, 46, 47).
- [10] Y. O. Dudin and A. Kuzmich, *Strongly interacting Rydberg excitations of a cold atomic gas*, *Science* **336** (2012) 887 (cit. on pp. 1, 41, 55).
- [11] D. P. Ornelas-Huerta et al., *On-demand indistinguishable single photons from an efficient and pure source based on a Rydberg ensemble*, *Optica* **7** (2020) 813 (cit. on p. 1).

- [12] H. Gorniaczyk, C. Tresp, J. Schmidt, H. Fedder and S. Hofferberth, *Single-Photon Transistor Mediated by Interstate Rydberg Interactions*, *Physical Review Letters* **113** (2014), ISSN: 1079-7114, URL: <http://dx.doi.org/10.1103/PhysRevLett.113.053601> (cit. on p. 1).
- [13] C. Tresp et al., *Single-Photon Absorber Based on Strongly Interacting Rydberg Atoms*, *Physical Review Letters* **117** (2016), ISSN: 1079-7114, URL: <http://dx.doi.org/10.1103/PhysRevLett.117.223001> (cit. on p. 1).
- [14] M. Fleischhauer and M. D. Lukin, *Dark-State Polaritons in Electromagnetically Induced Transparency*, *Physical Review Letters* **84** (2000) 5094, ISSN: 1079-7114, URL: <http://dx.doi.org/10.1103/PhysRevLett.84.5094> (cit. on pp. 1, 46, 47).
- [15] M. Saffman, T. G. Walker and K. Mølmer, *Quantum information with Rydberg atoms*, *Rev. Mod. Phys.* **82** (2010) 2313, arXiv: 0909.4777 [quant-ph] (cit. on pp. 1, 9, 41, 57, 62).
- [16] J. D. Pritchard et al., *Cooperative Atom-Light Interaction in a Blockaded Rydberg Ensemble*, *Physical Review Letters* **105** (2010), ISSN: 1079-7114, URL: <http://dx.doi.org/10.1103/PhysRevLett.105.193603> (cit. on p. 1).
- [17] T. Peyronel et al., *Quantum nonlinear optics with single photons enabled by strongly interacting atoms*, *Nature* **488** (2012) 57 (cit. on pp. 1, 63).
- [18] A. P. Burgers et al., *Controlling Rydberg Excitations Using Ion-Core Transitions in Alkaline-Earth Atom-Tweezer Arrays*, *PRX Quantum* **3** (2022), ISSN: 2691-3399, URL: <http://dx.doi.org/10.1103/PRXQuantum.3.020326> (cit. on p. 2).
- [19] K.-L. Pham, T. F. Gallagher, P. Pillet, S. Lepoutre and P. Cheinet, *Coherent Light Shift on Alkaline-Earth Rydberg Atoms from Isolated Core Excitation without Autoionization*, *PRX Quantum* **3** (2 2022) 020327, URL: <https://link.aps.org/doi/10.1103/PRXQuantum.3.020327> (cit. on p. 2).
- [20] X. Wang et al., *Two-color ytterbium magneto-optical trap in a compact dual-chamber setup*, *Phys. Rev. Appl.* **23** (1 2025) 014004, URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.23.014004> (cit. on pp. 2, 4).
- [21] A. V. Gorshkov, J. Otterbach, M. Fleischhauer, T. Pohl and M. D. Lukin, *Photon-Photon Interactions via Rydberg Blockade*, *Physical Review Letters* **107** (2011), ISSN: 1079-7114, URL: <http://dx.doi.org/10.1103/PhysRevLett.107.133602> (cit. on p. 2).
- [22] S. Weber et al., *Calculation of Rydberg interaction potentials*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **50** (2017) 133001, ISSN: 1361-6455, URL: <http://dx.doi.org/10.1088/1361-6455/aa743a> (cit. on pp. 2, 41, 43, 44, 72, 75).
- [23] J. Mögerle et al., *Accurate Modeling of Rydberg Atoms and Their Interactions: Theory and Implementation in PairInteraction*, (2026), arXiv: 2605.14993 [physics.atom-ph], URL: <https://arxiv.org/abs/2605.14993> (cit. on pp. 2, 41, 43, 44, 72, 75).

- [24] K. Kowalski et al., *Magneto-optical Trap: Fundamentals and Realization*, *Computational Methods in Science and Technology Special Issue* (2010) 115 (cit. on p. 4).
- [25] F. Pausewang, *Preparation and Detection of Ytterbium Rydberg Atoms and Molecules*, Master's thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2025, URL: <https://www.iap.uni-bonn.de/nqo/en/publications/publications-1/2025-florian-pausewang-master.pdf> (cit. on pp. 5, 51, 68).
- [26] E. J. Huland, *Precise optical trap alignment using piezo-motorized mirrors*, Bachelor's thesis: Institut für Angewandte Physik, Rheinische Friedrich-Wilhelms-Universität Bonn, 2026 (cit. on pp. 6, 8).
- [27] H. J. Metcalf and P. van der Straten, *Laser cooling and trapping of atoms*, *J. Opt. Soc. Am. B* **20** (2003) 887, URL: <https://opg.optica.org/josab/abstract.cfm?URI=josab-20-5-887> (cit. on p. 11).
- [28] C. J. Foot, *Atomic Physics*, Oxford Master Series in Physics, Oxford University Press, 2005 (cit. on pp. 11, 30).
- [29] Gooch & Housego, *AOMO 3200-129 Acousto-Optic Modulator, Datasheet*, 97-02197-01, 2013, URL: https://www.eqphotonics.de/wp-content/uploads/GH_AOM_3200-129_97-02197-01.pdf (visited on 28/09/2026) (cit. on p. 19).
- [30] R. Grimm, M. Weidemüller and Y. B. Ovchinnikov, *Optical Dipole Traps for Neutral Atoms*, *Advances in Atomic, Molecular and Optical Physics* **42** (2000) 95, arXiv: [physics/9902072](https://arxiv.org/abs/physics/9902072) [[physics.atom-ph](https://arxiv.org/abs/physics/9902072)] (cit. on pp. 19, 21, 30).
- [31] T. Pyragius, *Developing and building an absorption imaging system for Ultracold Atoms*, (2012), arXiv: [1209.3408](https://arxiv.org/abs/1209.3408) [[physics.ins-det](https://arxiv.org/abs/1209.3408)], URL: <https://arxiv.org/abs/1209.3408> (cit. on p. 20).
- [32] T. M. Brzozowski, M. Maczynska, M. Zawada, J. Zachorowski and W. Gawlik, *Time-of-flight measurement of the temperature of cold atoms for short trap-probe beam distances*, *Journal of Optics B: Quantum and Semiclassical Optics* **4** (2002) 62, URL: <https://doi.org/10.1088/1464-4266/4/1/310> (cit. on pp. 20, 53).
- [33] P. Bienias et al., *Photon propagation through dissipative Rydberg media at large input rates*, *Physical Review Research* **2** (2020), ISSN: 2643-1564, URL: <http://dx.doi.org/10.1103/PhysRevResearch.2.033049> (cit. on p. 24).
- [34] B. C. Das et al., *Long-range blockade between counter-propagating photons*, *Optica* **13** (2026) 1056, ISSN: 2334-2536, URL: <http://dx.doi.org/10.1364/OPTICA.586318> (cit. on p. 24).
- [35] T.-C. Poon and T. Kim, *Engineering Optics with MATLAB*, 2nd, World Scientific, 2018, ISBN: 978-981-3100-00-8 (cit. on p. 24).
- [36] J. R. Gilbert, C. P. Roberts and J. L. Roberts, *Near-resonant light propagation in an absorptive spatially anisotropic ultracold gas*, *Journal of the Optical Society of America B* **35** (2018) 718 (cit. on p. 24).
- [37] J. W. Goodman, *Introduction to Fourier optics*, vol. 1, 2005 (cit. on p. 26).

- [38] S. A. Self, *Focusing of spherical Gaussian beams*, *Appl. Opt.* **22** (1983) 658,
URL: <https://opg.optica.org/ao/abstract.cfm?URI=ao-22-5-658> (cit. on p. 26).
- [39] D. Maxwell et al., *Storage and Control of Optical Photons Using Rydberg Polaritons*, *Physical Review Letters* **110** (2013), ISSN: 1079-7114,
URL: <http://dx.doi.org/10.1103/PhysRevLett.110.103001> (cit. on pp. 41, 55).
- [40] I. I. Beterov, I. I. Ryabtsev, D. B. Tretyakov and V. M. Entin,
Quasiclassical calculations of blackbody-radiation-induced depopulation rates and effective lifetimes of Rydberg nS , nP , and nD alkali-metal atoms with $n \leq 80$,
Phys. Rev. A **79** (2009) 052504, Erratum: *Phys. Rev. A* **80**, 059902 (2009),
arXiv: [0810.0339](https://arxiv.org/abs/0810.0339) [[physics.atom-ph](https://arxiv.org/archive/physics)] (cit. on p. 42).
- [41] I. Mirgorodskiy et al.,
Electromagnetically induced transparency of ultra-long-range Rydberg molecules,
Phys. Rev. A **96** (1 2017) 011402(R),
URL: <https://link.aps.org/doi/10.1103/PhysRevA.96.011402> (cit. on pp. 46, 49, 60).
- [42] S. Sevinçli, N. Henkel, C. Ates and T. Pohl, *Nonlocal Nonlinear Optics in Cold Rydberg Gases*,
Phys. Rev. Lett. **107** (2011) 153001,
URL: <https://doi.org/10.1103/PhysRevLett.107.153001> (cit. on p. 46).
- [43] M. Fleischhauer, A. Imamoglu and J. P. Marangos,
Electromagnetically induced transparency: Optics in coherent media,
Rev. Mod. Phys. **77** (2 2005) 633,
URL: <https://link.aps.org/doi/10.1103/RevModPhys.77.633> (cit. on pp. 46, 47).
- [44] C. Mewes and M. Fleischhauer, *Decoherence in collective quantum memories for photons*,
Physical Review A **72** (2005), ISSN: 1094-1622,
URL: <http://dx.doi.org/10.1103/PhysRevA.72.022327> (cit. on p. 49).
- [45] B. Zhao et al., *A millisecond quantum memory for scalable quantum networks*,
Nature Physics **5** (2008) 95 (cit. on p. 52).
- [46] S. D. Jenkins, T. Zhang and T. A. B. Kennedy,
Motional dephasing of atomic clock spin waves in an optical lattice,
Journal of Physics B: Atomic, Molecular and Optical Physics **45** (2012) 124005,
URL: <https://doi.org/10.1088/0953-4075/45/12/124005> (cit. on pp. 53, 54).
- [47] T. Niederprüm, O. Thomas, T. Manthey, T. M. Weber and H. Ott,
Giant Cross Section for Molecular Ion Formation in Ultracold Rydberg Gases,
Phys. Rev. Lett. **115** (1 2015) 013003,
URL: <https://link.aps.org/doi/10.1103/PhysRevLett.115.013003>
(cit. on pp. 55, 56).
- [48] M. Schlagmüller et al., *Ultracold Chemical Reactions of a Single Rydberg Atom in a Dense Gas*,
Phys. Rev. X **6** (3 2016) 031020,
URL: <https://link.aps.org/doi/10.1103/PhysRevX.6.031020> (cit. on pp. 55, 56, 58).

- [49] I. Beigman and V. Lebedev,
Collision theory of Rydberg atoms with neutral and charged particles,
Physics Reports **250** (1995) 95, ISSN: 0370-1573,
URL: <https://www.sciencedirect.com/science/article/pii/037015739500074Q>
(cit. on pp. 55, 56, 58).
- [50] C. Chin, R. Grimm, P. Julienne and E. Tiesinga, *Feshbach resonances in ultracold gases*,
Reviews of Modern Physics **82** (2010) 1225, ISSN: 1539-0756,
URL: <http://dx.doi.org/10.1103/RevModPhys.82.1225> (cit. on p. 58).
- [51] M. Kitagawa et al., *Two-color photoassociation spectroscopy of ytterbium atoms and the precise determinations of s-wave scattering lengths*, *Physical Review A* **77** (2008), ISSN: 1094-1622,
URL: <http://dx.doi.org/10.1103/PhysRevA.77.012719> (cit. on p. 58).
- [52] T. Legrand et al.,
Revealing Electron-Ytterbium Interactions through Rydberg Molecular Spectroscopy,
Physical Review X **16** (2026), ISSN: 2160-3308,
URL: <http://dx.doi.org/10.1103/tjv3-485v> (cit. on pp. 59–61).
- [53] M. Fox, *Quantum Optics: An Introduction*, Oxford: Oxford University Press, 2006,
ISBN: 9780198566724 (cit. on p. 63).