Room Temperature Single-Photon Sources and Atomic Quantum Memories for Broadband Quantum Networks

Inauguraldissertation

zur
Erlangung der Würde eines Doktors der Philosophie
vorgelegt der
Philosophisch-Naturwissenschaftlichen Fakultät
der Universität Basel

von

Roberto Mottola

2024

Genehmigt von der Philosophisch-Naturwissenschaftlichen Fakultät
auf Antrag von
Erstbetreuer: Prof. Dr. Philipp Treutlein
Zweitbetreuer: Prof. Dr. Richard Warburton
externer Experte: Dr. Mikael Afzelius

Basel, den 21. Februar 2023

Prof. Dr. Marcel Mayor
Dekan

Abstract

Quantum networks are envisioned to overcome current limitations in quantum communication and computation. The building blocks that enable the realization of such networks are quantum memories and single-photon sources. A promising path to realize these networks is to make use of heterogeneous interconnects. By interfacing quantum memories based on atomic ensembles with solid-state sources, the best of two worlds could be exploited. However, achieving compatibility between different systems, in order to realize a hybrid quantum network node, has remained an open challenge.

In this thesis I report on broadband quantum memories implemented in warm rubidium vapor and on a compatible single-photon source based on non-degenerate cavity-enhanced spontaneous parametric downconversion. The goal is to implement an elementary interconnect where the experimental complexity is kept low by operating all components at or above room temperature. Choosing a monolithic cavity design, the source is inherently robust and reaches high efficiencies. It generates heralded single photons with hundreds of MHz bandwidth and reaches high heralding efficiencies of $\geq 40\,\%$, measured after coupling into a single-mode optical fiber.

The memories presented here are based on electromagnetically induced transparency (EIT) in a lambda-level scheme in warm rubidium vapor. To suppress the read-out noise, which is a limiting factor in common ground-state memory implementations, two separate approaches are followed. In the first one, the atomic structure is modified by applying a tesla-order magnetic field and working in the hyperfine Paschen-Back regime. This results in large splittings between the energy levels, which allows us to optically address individual sublevels in the warm vapor. A spectroscopic study of EIT and optical pumping in this regime is presented. Our proof-of-principle implementation of a quantum memory in a miniaturized vapor cell has delivered promising results. This setup was capable of storing and retrieving weak coherent pulses attenuated to the single-photon level, yielding an end-to-end efficiency of $\eta_{e2e} \approx 3\,\%$ and a SNR of up to 7.9(8). The second approach to suppressing read-out noise relies on exploiting polarization selection rules in a Zeeman-pumped vapor. This memory implementation was successfully interfaced with 370 MHz-broad single photons from the heralded downconversion source. The stored photons maintained their nonclassical signature after retrieval, yielding a $g^{(2)}(0) = 0.177(23)$. This constitutes the first demonstration of single photon storage and retrieval from the ground state of a warm atomic vapor.

The developed platform operates in a technologically relevant regime for future experiments, paving the way for the exploration of promising quantum-network protocols at high bandwidth.

Contents

Abbreviations

AOM	Acousto-optic modulator
AR	Anti-reflection (coating)
ASE	Amplified spontaneous emission
AT	Autler-Townes
AWG	Arbitrary waveform generator
CPT	Coherent population trapping
CW	Continuous wave
DDG	Digital delay generator
DFB	Distributed feedback (laser)
DFG	Difference frequency generation
ECDL	External cavity diode laser
EIT	Electromagnetically induced transparency
EOM	Electro-optic modulator
FSR	Free spectral range
FWM	Four-wave mixing
OD	Optical depth
OI	Optical isolator
HBT	Hanbury Brown and Twiss
HPB	Hyperfine Paschen-Back (regime)
IF	Interference filter
ND	Neutral density (filter)
NIR	Near-infrared
OPO	Optical parametric oscillator
PD	Photodiode
PM	Polarization maintaining (fiber)
ppKTP	Periodically poled potassium titanyl phosphate
QFC	Quantum frequency conversion
SM	Single mode (fiber)
SNR	Signal-to-noise ratio
SNSPD	Superconducting nanowire single-photon detector
SOA	Semiconductor optical amplifier
SPAD	Single-photon avalanche detector

SPDC	Spontaneous parametric downconversion
TA	Tapered amplifier
TCSPC	Time correlated single photon counting
TEC	Thermoelectric cooler
WCP	Weak coherent pulse

1 Introduction

The internet had a revolutionary impact on our society. A similar breakthrough in quantum information science is envisioned with the realization of quantum networks. Such networks are composed of nodes that store and process quantum information and that are linked via quantum channels. Single photons are, due to their weak interaction with the environment and their quantum nature even at room temperature, the leading candidate to exchange quantum information from one node to another. They can be used to directly transfer qubits, enable sharing the resource of entanglement, or synchronize probabilistic operations. Optical interconnects that reversibly map quantum states between matter and light are the fundamental building blocks of quantum networks [1]. For this endeavor, quantum memories [2, 3, 4, 5, 6] and single-photon sources (SPSs) [7] constitute enabling technologies.

The roadmaps for a quantum internet [8, 9] envision networks capable of quantum communication [10] made secure through quantum key distribution (QKD) [11, 12]; quantum information processing (QIP) [13] such as distributed [14, 15, 16, 17] or blind quantum computing [18]; and quantum metrology [19], including distributed quantum sensing [20], clock synchronization [21], and even enhancing the resolution of interferometric astronomical telescopes [22]. As with any transformative technology, it is difficult to predict all uses.

Recently, the first elementary quantum networks consisting of three nodes have been realized. The entanglement of the nodes via three-photon interference was demonstrated in cold Rb ensembles [23], while a three-node entanglement-based network in nitrogen-vacancy centers [24] successfully performed quantum teleportation between non-neighboring nodes [25]. There is a large variety of different systems to choose from to implement the necessary network components, each with its strengths and weaknesses. No platform so far was able to satisfy all requirements simultaneously. Most likely, the requirements of different applications will be best fulfilled by different platforms, similarly to classical memories, which apply specific technologies depending on the use-case. For harnessing the advantages of different systems, heterogeneous approaches where interconnects of different technologies are interfaced are being researched as well. Successful demonstrations of hybrid interconnects have been achieved, for example by entangling two different rare-earth ion quantum memories working at different wavelengths through spontaneous parametric downconversion (SPDC) photons [26]. Furthermore, a fully heterogeneous interconnect has been demonstrated between a cold atomic ensemble and a rare-earth doped crystal mediated by quantum frequency conversion (QFC) [27]. First successes of heterogeneous interfaces have also been recently reported on with warm atomic memories [28, 29].

A prominent protocol for the robust distribution of quantum information over networks is the quantum repeater [30]. For connecting the various network nodes, free space and optical fibers are suitable transmission channels for photons. Free-space links between satellites and ground stations have been established with single photons over distances of up to 1400 km [31, 32], while in ultralow-loss optical fibers, distances of ~ 830 km have been successfully bridged. These demonstrations, however, have reached or are close to reaching the limits of what is possible with direct links, due to fundamental bounds on the transmission rates [33, 34]. In classical communication, transmission loss can be compensated with signal amplification or repetition, and thus does not limit the maximum distance between network links. However, in quantum communication the no-cloning theorem [35] rules out the repetition or signal amplification to overcome link imperfections and lossy channels.

Quantum teleportation [36] can be used to overcome these limitations. While there is no fundamental distance limit for teleportation, the entanglement distribution needed for this protocol will, however, be subject to the limit set by the maximum rate at which distant nodes can be entangled over lossy channels [37]. In order to generate long-distance entanglement the quantum repeater protocol was proposed: instead of relying on direct transmission between two nodes, the distance is split into elementary links. First, entanglement is created between adjacent nodes separated by a distance that can be bridged by direct transmission. Once all pairs of neighboring links have succeeded with establishing entanglement, it can be distributed over longer distances by entanglement swapping [38] until the two outermost locations share entanglement. The process of distributing entanglement is probabilistic, thus succeeding at different times at different nodes. For the entanglement swapping to work, however, it needs to succeed simultaneously across all elementary links. The entanglement generation can be synchronized with quantum memories, i.e. the swap operation is delayed until all adjacent links actually share entanglement.

An influential proposal for the realization of a quantum repeater was put forward in 2001 by Duan, Lukin, Cirac, and Zoller (DLCZ) [39]. This protocol uses atomic ensembles as quantum memories, linear optics, and single-photon counting to perform all necessary operations. The authors motivate the choice of an ensemble over a single quantum system (e.g. a single atom) because of its simpler realization and a greatly improved generation rate. The latter is achieved thanks to the collective enhancement due to the large number of atoms. Various theoretical extensions and several refinements followed for this protocol, using photon-pair sources and multimode memories [40, 41], as well as several proof-of-principle experimental implementations [42, 27].

Sangouard *et al.* [43] have proposed an efficient quantum repeater architecture based on SPSs and quantum memories at each node, which is particularly relevant for the work in this thesis. By eliminating the intrinsic two-photon-pair emission probability, which limits the performance of the DLCZ scheme, this protocol can significantly enhance the achievable entanglement distribution rates. Each network node has an SPS that ideally creates one photon when excited. This photon is sent onto an unbalanced beam splitter with the quantum memory at one output,

while the other mode is coupled into a fiber leading to a central station. At this central station the modes coming from two separate nodes are combined onto a 50:50 beam splitter. The detection of a single-photon after the splitter erases the which-way information creating entanglement between the two memories, which are in a superposition of being empty or having stored a single photon.

Quantum memories and SPSs constitute key elements of quantum repeaters. They are necessary for the generation of single photons and for storing and retrieving them on-demand. Key requirements for the SPSs and memories are high emission rates, long storage times with high efficiencies, and preferably multi-mode capability. For implementing a scalable network it would be ideal if memories and sources can be operated at room temperature, drastically reducing the complexity of the infrastructure needed at each network node.

Quantum networks do not need to bridge long distances to be interesting. Smaller-scale, local quantum networks can also be beneficial, without the necessity of telecom wavelength photons. For example, a viable application would be quantum computing on a network. The scalability of quantum computers is not a trivial task [44], for various reasons depending on implementation that range from architecture designs, to space requirements on the chips, and the dissipation of heat induced by the read-out from the qubits in the dilution fridge. A local network could connect several such 'quantum processors' for scaling the computational resources enabling distributed quantum computing, which would not require (or even benefit from) large distances between the nodes.

The applications for quantum memories and SPSs are not limited to quantum networks. Optical QIP as linear optical quantum computation (LOQC), for example, requires only linear optical elements, namely beam splitters and phase shifters, SPSs, and single-photon detectors [45]. The nonlinearity necessary for universal computing is obtained probabilistically through photon measurements. With a quantum memory, each operation is to be repeated until successful, with the output subsequently stored until needed. Measurement based quantum computing (MBQC) [46], which requires the scalable generation of cluster states, would also benefit from using memories [47]. Quantum memories can strongly increase the success probability of multi-photon states through synchronization. It is advantageous to place memories downstream of heralded photon sources and fill the memories as soon as a photon is emitted. Once all memories are full and the last photon is generated, they can all be read out at once. The generation of such states with N coincident photons can find direct application in boson sampling experiments [48]. In [49] the authors identified the time-bandwidth product of a quantum memory as the most important figure of merit for synchronization tasks. For this type of applications memory and detector efficiencies and fidelity are other important figures of merit. Beyond synchronization tasks, some implementations of quantum memories can be used as versatile QIP devices, e.g. by exploiting the selectivity of specific time modes [50] or by engineering the optical state by multi-pulse addressing [51].

In this thesis, we investigate quantum memories implemented in hot alkali vapor

[3, 6]. They constitute an attractive system in which photons with bandwidths of several hundreds of MHz up to GHz – significantly larger than the natural linewidth of the atomic transitions – can be stored. The storage is induced by applying control pulses with high Rabi frequencies that map an incoming broadband photon onto a shared atomic excitation in the long-lived ground state. Upon sending a second control pulse this spin wave is converted back into a photonic excitation. The experimental simplicity of hot alkali vapor memories is promising for scalability. MEMS (microelectromechanical systems) technology has led to the successful miniaturization of vapor cells for quantum sensors, especially atomic clocks [52], which, if paired with low-power lasers, could be integrated in portable devices. These properties are promising when considering scalable quantum networks, the ability for spatial-multiplexing, and potentially even enabling satellites-borne applications [53]. The wavelengths of alkali metals, furthermore, are convenient due to the availability of efficient single-photon detectors and high-power lasers. Successfully interfacing a ground-state atomic vapor quantum memory with single photons, however, still remained an open challenge. The generated read-out noise constitutes a limiting factor in preserving the non-classical properties of the retrieval [54, 55, 56]. Furthermore, for interfacing, it is necessary that the SPS and the memory are bandwidth- and wavelength-matched. For these reasons, most atomic vapor memory experiments use laser pulses attenuated to the single-photon level [57, 58, 56, 59].

The goal of this thesis is to explore the feasibility of warm alkali vapor based quantum memories and SPSs as building blocks for elementary quantum network nodes. A hybrid quantum-node approach is followed, where SPS and memory are realized in different technologies. On the one hand, we want to exploit the long coherence times of atomic ground states for quantum memories, and on the other hand, we take advantage of the excellent SPS qualities of solid-state sources. We focus on a scalable, high-bandwidth implementation at near-infrared (NIR) wavelengths, for which efficient sources and memories can be developed, while efficient detectors are commercially available. One of these memories is even implemented in a miniaturized vapor cell. The working principle of our quantum memories is based on electromagnetically induced transparency in a lambda-scheme. However, we operate the memories with a small detuning from resonance, on the order of the Doppler-broadening, for reducing the read-out noise. We investigate two opposite approaches to select and engineer appropriate three-level systems in the atoms, with the aim to reduce spurious processes, particularly the generation of read-out noise and parasitic absorption: 1) By applying a tesla-order magnetic field we take advantage of the Zeeman effect and enter the hyperfine Paschen-Back (HPB) regime. The induced energy splitting between the levels is significantly larger than the optical linewidth, allowing us to address single transitions, which usually is only possible in cold atoms. Before the memory was implemented, this regime was studied by performing spectroscopic analysis of EIT and optical pumping. 2) In a zero-field scenario we Zeeman-pump the atoms and exploit polarization selection rules to suppress noise channels that so far have been the limiting factor. We operate our memories in a high-bandwidth regime (compared to the natural linewidth of alkali metals) of several hundreds of MHz up to almost one GHz, making them especially interesting for applications that

require fast operation, such as local networks, synchronization tasks and QIP.

This operating regime makes these ensemble-based memories interesting for interfaces with photons emitted by GaAs quantum dots [60, 61, 62, 63]. While currently neither system was sufficiently mature for interfacing attempts within the scope of this thesis, it is promising that such an interconnect could be realized in the near future. We interfaced our memories with single photons generated by SPDC. This nonlinear process that probabilistically generates photon pairs, has the great advantage of working at room temperature. Detecting one photon of a pair announces the presence of the second, constituting a heralded single-photon source. Cavity enhancement allows us to tailor the emission bandwidth of this nonlinear process from nm down to hundreds of MHz to match the acceptance bandwidth of the vapor memory. Choosing a monolithic design, where the downconversion crystal functions simultaneously as the cavity, is intrinsically robust and helps to significantly reduce the losses and thus to obtain a high heralding efficiency. We combined this heralded single-photon source with the Zeeman-pumped memory, demonstrating for the first time the successful storage and retrieval of single photons in a broadband ground-state memory implemented in a hot vapor.

Some of the results presented in this thesis are also published in the following papers:

- R. Mottola, G. Buser, C. Müller, T. Kroh, A. Ahlrichs, S. Ramelow, O. Benson, P. Treutlein, and J. Wolters, *An efficient, tunable, and robust source of narrowband photon pairs at the ^{87}Rb D1 line*, Optics Express **28**, 3159 (2020)
- G. Buser, R. Mottola, B. Cotting, J. Wolters, and P. Treutlein, *Single-photon storage in a ground-state vapor cell quantum memory*, PRX Quantum **3**, 020349 (2022)
- R. Mottola, G. Buser, and P. Treutlein, *Optical memory in a microfabricated rubidium vapor cell*, Physical Review Letters **131**, 260801 (2023)
- R. Mottola, G. Buser, and P. Treutlein, *Electromagnetically induced transparency and optical pumping in the hyperfine Paschen-Back regime*, Physical Review A **108**, 062820 (2023)

Structure of the thesis This thesis is organized as follows. Chapter 2 establishes the basic concepts of solid-state single-photon sources and atomic vapor memories. The focus is set on the systems developed in the frame of this thesis, namely heralded downconversion sources and EIT-based quantum memories. The chapter concludes with a brief description of the numerical simulation of the storage and retrieval process, implemented to analyze limitations and guide future developments of the memory experiments.

In Chapter 3 the monolithic SPDC source is presented. The various design iterations of the source and their respective characterization are described. The causes of the

observed crosstalk between the source and the memories are investigated and the adaptations necessary for a successful interfacing of the two systems are explained.

The first part of Chapter 4 describes the HPB regime. It starts with a brief theoretical introduction to the Paschen-Back regime and to the light-matter interaction at high magnetic fields. It continues with spectroscopy and EIT measurements as a preparation for the quantum memory implementation. The final part of the chapter presents first attempts to polarize the nuclear spin of the atoms in this regime.

Chapter 5 is devoted to the HPB memory. The memory scheme for the high magnetic field scenario is introduced and it is explained how a near-ideal three-level system is isolated. A thorough description of the experimental setup is followed by storage and retrieval experiments with weak coherent pulses (WCPs). The chapter ends with first interfacing experiments with the SPDC source, discussing possibilities for future improvements.

Chapter 6 reviews the second approach we pursue to engineer a clean lambda system: Zeeman-pumping the atomic ensemble. The experimental results demonstrating the first ever non-classical storage and retrieval in a broadband ground-state atomic-vapor memory are presented. The chapter closes with some insights obtained by the numerical simulation.

The thesis is concluded in Chapter 7 with a recapitulation of the achieved experimental results. Open challenges for the presented systems and future goals towards quantum network applications are summarized.

2 About Single-Photon Sources and Quantum Memories

In the introduction we saw that single-photon sources (SPS) and quantum memories can be considered enabling technologies for quantum communication and networking applications. In this chapter I will give a short overview of SPSs and memories and their respective figures of merit. Due to their breadth, this is not intended to be a complete review of these two fields. Rather, I will focus on the systems I implemented experimentally and those closely related, or ones that might be interesting for the direct future development of our experiments.

2.1 Single-Photon Sources

An ideal SPS [7] can be thought of as an apparatus that can emit a single photon 'on-demand', meaning at an arbitrary point in time chosen by the experimenter. Upon being triggered, this device should emit exactly one photon with certainty. Subsequently emitted photons should, furthermore, be indistinguishable. In addition, the repetition rate of such an ideal source should only be limited by the temporal duration of the emitted photons. Real-world implementations of SPSs, however, deviate from the desired properties and trade-offs need to be made, always with a specific experimental application in mind.

The development of such sources is a flourishing field of research, with many different physical platforms being investigated. In general, a distinction between deterministic and probabilistic SPSs is warranted. Deterministic systems emit a single photon on-demand. Mostly, systems composed of single emitters are considered for this type of source implementation, such as quantum dots [67, 60], diamond color centers [67, 68, 69], single atoms [70], single ions [71], and single molecules [72]. The underlying photon generation process in these diverse systems is similar. The operating principles rely on an external control, e.g. a laser pulse, being applied to trigger the emission of the single photon. The systems are excited by this control and will emit a photon upon relaxation to a lower energy state. The emission of single emitter sources is thus defined by the transition linewidth and wavelength. Optical cavities are often used to modify the emission characteristics. On the other hand, probabilistic SPSs rely upon the generation of a correlated photon pair through spontaneous processes. One of the two photons of the pair can be detected to 'herald' the presence of the other photon. The ability to herald an emitted single photon makes such sources interesting for quantum information applications, even

though they are inherently of probabilistic nature. Common implementations are SPDC in $\chi^{(2)}$ non-linear crystals [73, 74, 75, 76] and waveguides [77, 78], and spontaneous four-wave mixing (FWM) in $\chi^{(3)}$ non-linear fibers [79, 80, 81], waveguides [82, 83, 84], but also in atomic ensembles [85, 86].

Before going into detail describing selected SPS platforms, I will discuss the relevant figures of merit necessary to compare the various systems and to establish a common language.

Second-Order Correlations and 'State Accuracy' One way to assess the quality of the output of a SPS is to determine the overlap of the reconstructed density matrix of an emitted photon with that of a Fock state with $n = 1$. This, however, requires quantum-state tomography [87], which is an experimentally arduous and time-consuming task. Conveniently, metrics like the normalized Glauber second-order correlation, sometimes also referred to as the "degree of coherence", can be used to quantify the multiphoton component of a state. I refer the interested reader to one of many quantum optics textbooks, such as [88] or [89], for its derivation.

The second-order correlation is a function of the delay time τ between the photon detections and is given by

$$g^{(2)}(\tau) = \frac{\langle \hat{E}^{(-)}(t)\, \hat{E}^{(-)}(t+\tau)\, \hat{E}^{(+)}(t+\tau)\, \hat{E}^{(+)}(t) \rangle}{\langle \hat{E}^{(-)}(t)\, \hat{E}^{(+)}(t) \rangle \langle \hat{E}^{(-)}(t+\tau)\, \hat{E}^{(+)}(t+\tau) \rangle} \,. \tag{2.1}$$

Here $\hat{E}^{(+)}$ is the positive frequency component of the electromagnetic field proportional to the photon annihilation operator and thus describing absorption, while $\hat{E}^{(-)}$ is its Hermitian conjugate. The $g^{(2)}(\tau)$ can be interpreted as the probability of a second detection a time delay τ after a first photon was detected.

For a single-mode field this quantity can be simplified to

$$\begin{aligned} g^{(2)}(\tau) &= \frac{\langle \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} \rangle}{\langle \hat{a}^\dagger \hat{a} \rangle^2} \\ &= 1 + \frac{\langle (\Delta \hat{n})^2 \rangle - \langle \hat{n} \rangle}{\langle \hat{n} \rangle^2} \end{aligned} \tag{2.2}$$

and becomes independent of τ. For a Fock (or number) state the expression above can be further simplified to $g^{(2)}(\tau) = 1 - 1/n$. For an ideal single-photon state, corresponding to the $n = 1$ Fock state, we get $g^{(2)}(0) = 0$.

The $g^{(2)}(\tau)$ can take classically forbidden values, thus yielding two indicators for non-classical light: one is sub-Poissonian statistics identified by $g^{(2)}(\tau) < 1$, and the other is photon antibunching, described by $g^{(2)}(0) < g^{(2)}(\tau)$. The latter phenomenon characterizes a source for which it is more likely that successive photons are emitted with long time separation, rather than at short time intervals. For a single-mode thermal state (or chaotic light) $g^{(2)}(0) = 2$ is expected, which describes the photon bunching effect, i.e. the photons are more likely to be arrive simultaneously and

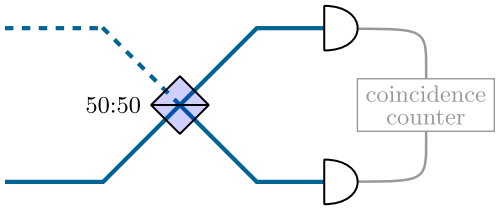

Figure 2.1: Sketch of a Hanbury Brown Twiss interferometer. The output of a SPS is sent onto a balanced beam splitter (solid line). A detector is placed after each output and the coincidences are counted with a time-to-digital converter. The setup for a HOM interferometer is the same except that two successively emitted photons are each sent to a different input of the beam splitter.

induce coincidences on the detectors. For a coherent state, as laser light is commonly modeled, a stream of photons with no correlation is generated, thus resulting in $g^{(2)}(\tau) = g^{(2)}(0) = 1$. With an increasing number of modes, a multimode state is expected to reach the same value.

In some communities, especially the one studying solid-state emitters, the term 'single-photon purity', being defined as $1 - g^{(2)}(0)$, is adopted to describe the quality of the output of a SPS. Since the term purity also holds other meanings, to avoid confusion, I prefer the terms '(single-photon) state accuracy' and 'single-photon nature' instead, as used in [7].

The single-photon state accuracy is one of the key metrics used for the experiments described in this thesis. It can be measured by directing the output of the SPS (or the retrieved signal of a quantum memory) to a Hanbury Brown Twiss (HBT) interferometer (see Fig. 2.1). The second-order (auto)correlation can be expressed as ratio of the photon detection probabilities [89]

$$g^{(2)} = \frac{p\,(d1 \wedge d2)}{p\,(d1)\,p\,(d2)}\,. \tag{2.3}$$

The numerator represents the probability of measuring a coincidence, while $p(d1)$ and $p(d2)$ describe the probability of a click on either detector. These probabilities can be easily estimated from the counts detected within a certain time frame.

For heralded photons, the second-order correlation conditioned on the detection of a herald photon is the measure for determining the single-photon nature of the output. This quantity depends on events where all three detectors (herald and two HBT) click [90].

Indistinguishability The photon indistinguishability can be seen as the overlap of the quantum states of two photons, which are either successively emitted by the same source or originate from different sources. This can be quantified by performing a two-photon interference measurement, known as Hong-Ou-Mandel (HOM) interference [91]. The two photons are each sent onto one input of a 50:50 beam

splitter, while a single-photon detector is placed behind each output. Two indistinguishable photons will perfectly interfere on the beam splitter, resulting in no coincidences between the detector clicks. This can be shown by rigorously calculating the output wave function of a balanced beam splitter for a single-mode input [92] or can be intuited from the Bosonic properties of photons. They require the global wave function for the two photons to be symmetric. Thus, if the photons are indistinguishable in every other degree of freedom, they must exit the beam splitter through the same output to yield a spatially symmetric wave function as well [93]. Usually, the coincidences of a HOM measurement are recorded as a function of a variable time delay. The indistinguishability is then quantifiable through the HOM visibility $\mathcal{V}$ defined by $\mathcal{V} = (C_{max} - C_{min})/C_{max}$ with C being the maximum and minimum detected coincidence rate, respectively. We will see later when discussing quantum memories that the photon indistinguishability can be seen as the analog to a memory's fidelity.

Efficiency The efficiency can be understood as the probability of a single photon to be generated into a specific mode. A distinction is necessary between deterministic and probabilistic sources. For deterministic SPSs the end-to-end efficiency is constituted by the product of the excitation probability times the extraction efficiency. The first one constitutes the probability that the physical system is excited upon application of a trigger, while the latter determines how efficiently the generated photon can be extracted from the medium into a specific mode of interest. Unfortunately, the definitions of 'extraction efficiency' found in the literature varies, and it is often defined at some arbitrarily chosen point, for example at the first lens as in [60]. In my opinion, such a definition is not appropriate as a specification of how application-suited a source is. For such a purpose, I prefer the definition of the likelihood to successfully extract the photon out of its environment into a single-mode fiber. Not only does this definition constitute an end-to-end efficiency for the source (when multiplied by the excitation probability), which is the relevant quantity for applications, but it also ensures that the emission is single mode. Furthermore, it is usually necessary to couple the output into a fiber for detection anyhow, and this way the source is directly ready for any interface that accepts a fiber input.

Since an isolated mode of a photon-pair source only exhibits single-photon characteristics if conditioned on the other mode, it is convenient to define the heralding efficiency for probabilistic sources. This is the probability of having a signal photon in the output fiber of the SPS after having detected a herald photon. It is important to point out the asymmetry arising from this conditioning. While losses in the signal arm have direct repercussions on the source's efficiency, losses in the herald arm only reduce the rate of conditioning.

The efficiency of an SPS is particularly important when considering interfaces with a quantum memory affected by read-out noise, as is the case for hot-vapor ground-state memories. In fact, every time the memory is triggered without a single photon coming from the source, the read-out noise induced by the memory's control is accumulated. A high noise contribution will contaminate the output correlations

and deteriorate the single-photon nature of the retrieval. It is therefore convenient to define a μ_1 parameter for a memory that states the necessary minimal efficiency of a source to reach a signal-to-noise ratio of one for said memory [94].

Generation Rate The generation rate, in some cases referred to as brightness, refers to the single photons that are emitted per unit of time by the source. In the case of deterministic SPSs, it is determined by the repetition rate of the exciting laser, however, it is limited by the lifetime of the excited state. For probabilistic sources, on the other hand, the best metric to describe the emission rate are the herald detections per second.

By our definition of an ideal SPS, a multimode output is not desired, since it would make the emitted photons distinguishable. If a device emits two separable modes, it might be considered as two individual SPSs. By defining the efficiency after coupling the output into a single-mode (SM) fiber, which acts as a spatial filter, we have already narrowed down the emission to a sole spatial mode. If a source emits different spectral modes they need to be filtered, e.g. with etalons. If the time-bandwidth product of the emission is Fourier limited the source only emits one spectral mode [95]. For heralded SPS the number of modes can be estimated by the peak height reached by the unconditioned autocorrelation (see Ch. 3).

All the figures of merit presented up to here describe how closely a source approaches an ideal SPS. If, however, the SPS is meant to be interfaced with another system, as for example in our case a quantum memory with the application of building quantum-network nodes, the compatibility of the two systems is also important. For such an application, wavelength and bandwidth matching are important for successfully interfacing the systems. While this appears trivial when implementing the source in the same physical system as the addressed memory, e.g. both SPS and memory are in an atomic system, it can be quite a challenging engineering task to match solid state sources with atomic memories.

2.1.1 Atomic Ensembles

The photons emitted by SPSs based on atomic ensembles are inherently compatible with systems of the same atomic species. The most prominent example of SPSs based on atomic ensembles relies on the Duan-Lukin-Cirac-Zoller (DLCZ) protocol proposed for quantum networks [39]. This protocol relies on atoms with a lambda-scheme of energy levels. After preparing the ensemble into one ground state, a short, off-resonant laser pulse induces Raman scattering. A Stokes photon announces the scattering event and the creation of a shared atomic excitation, which is then ready to be read out on demand. By applying a second control pulse on-demand, the source is triggered and the atomic excitation is converted into a photon. The source acts like the quantum memories presented later (see Section 2.2) except for the fact that the ground-state spin wave is generated spontaneously through scattering, as opposed to being deterministically read-in from an 'input' signal. The DLCZ read-only memory

is a heralded SPS that, however, can emit its photons with a variable delay, which is limited by the coherent excitation's lifetime. Implementations of this scheme can be found both in cold and hot atomic ensembles. Even a fully heterogeneous quantum network interconnect has been demonstrated using this protocol [27].

Heralded SPSs can be implemented in atomic ensembles as well. A recent implementation [86] is based on spontaneous FWM in a warm Rb ensemble, where GHz-broad photons with a HOM visibility of $> 80\,\%$ were generated at high rates with low noise.

2.1.2 Quantum Dots

Quantum dots are (potentially) deterministic solid-state emitters made of a low bandgap semiconductor material embedded in a high bandgap material. The most investigated type of quantum dots is composed of InAs or InGaAs in GaAs. This type of quantum dots excels with $g^{(2)}(0)$ values below $1\,\%$ [60]. At the same time, the photons emitted by the same quantum dot can reach indistinguishabilities above $99\,\%$ [96]. The operation rate is determined by the radiative lifetime of the emitter, which in some cavity-based approaches can reach repetition rates of $1 - 2\,\mathrm{GHz}$ [60]. However, their typical emission wavelengths are $\geq 900\,\mathrm{nm}$, making them incompatible with memories in Rb vapor.

During the growth process, the lattice mismatch between the utilized materials leads to the creation of small islands (in the order of tens of nm) of the low bandgap material: the quantum dots. Through this strong spatial confinement of a quantum dot, the resulting potential experienced by electrons and holes can be described as square potential with its characteristics discrete energy levels. In a semiconductor the occupied valence band and the unoccupied conduction band are separated by the bandgap. A photon can excite an electron into the conduction band, creating a hole. The electron-hole pair is called an 'exciton'. The ground state and the exciton constitute a two-level system, resembling that of an atom [97].

Due to the high refractive index of GaAs the probability of the emitted photons exiting the host material from one facet is only around $2\,\%$ [60]. In order to increase the extraction efficiency and couple the generated photons into one optical mode, various strategies to engineer the photonic environment have been adopted, such as coupling the quantum dots to micropillars, nanowire antennas, microcavities, and more.

The recent developments of GaAs/AlGaAs quantum dots are of particular interest for our purposes. In fact, these semiconductor emitters can be considered to be "artificial Rb atoms" [98]. The fabrication process includes etching nanoholes with Al droplets and successively filling them with GaAs. During fabrication, the emission wavelength can be coarsely tuned to reach about $780\,\mathrm{nm}$, matching the Rb D_2 line. The self-assembly fabrication process results in a spread of wavelengths of the quantum dots across a sample. A single emitter can be tuned by inducing strain with

a piezoelectric actuator for fine-tuning to a specific Rb transition line, with tuning ranges up to 1.15 THz, which is approximately three orders of magnitude larger than its linewidth [99]. To suppress charge noise (which initially was worse than in InGaAs dots), the dot is embedded in a n-i-p diode structure. This results in a blinking-free emission with transform-limited linewidths [61], while adding the ability to tune the emission wavelength through applied gate voltages. With these improvements, a high indistinguishability of photons emitted by independent quantum dots, reaching visibilities of 93 %, was measured without any kind of filtering [62]. So far, this was a longstanding problem, even for the 'standard' quantum dots, since the interference is subject to uncorrelated noise in the two separate semiconductor environments. Often, photons from different dots had to be filtered (spectrally or temporally) to make them indistinguishable [100].

It has been demonstrated that by exploiting spontaneous spin-flip Raman transitions designer photons can be generated in GaAs quantum dots [101]. The shape of the emitted photons is determined by the tailored temporal profile of the applied excitation pulse. Bandwidths as low as 200 MHz, constituting a reduction of almost an order of magnitude, were achieved, making this an interesting technique to reach bandwidth matching with other systems.

In some recent work [63], the on-demand generation of single photons with a probability of 57 % to have the photon at the end of the final fiber was demonstrated at a rate of 1 GHz using an InGaAs quantum dot coupled to an open microcavity. This is a huge advancement closing the compatibility gap between semiconductor quantum dots and atomic memories a little further.

The development currently in progress for this new type of GaAs quantum dots is promising. Some transfer of techniques used on their well established InGaAs counterparts is still necessary, though. It seems, however, that, e.g., adapting microcavities to these 'Rb-like' emitters should only be a matter of time. Furthermore, fabrication of such dots at 795 nm is being currently tested. These advances would allow us to consider quantum dots as sources to interface our vapor memories with. Richard Warburton's research group at the Universität Basel, with which we have an ongoing collaboration, is one of the pioneers in this field.

2.1.3 Spontaneous Parametric Downconversion

Parametric downconversion is the most common process used for the generation of correlated photon pairs. This generation process is experimentally uncomplicated and can be performed at room-temperature. Furthermore, achieving tunability over a wide range of wavelengths is possible [102]. These features made it attractive for us to test our memories with a downconversion source first. However, the low generation rates of the process can constitute a major problem for some applications. Low rates cannot be compensated by increasing pump power indefinitely, since this inevitably results in a deterioration of the state accuracy.

In the SPDC process a pump beam illuminates a material with a second-order nonlinear optical susceptibility $\chi^{(2)}$. Within the medium an incoming photon from the pump beam is spontaneously downconverted in a three-wave mixing process into two photons with lower frequency under conservation of energy and momentum:

$$\omega_{\text{pump}} = \omega_{\text{signal}} + \omega_{\text{idler}} \ , \qquad \Delta k = 0 \ , \tag{2.4}$$

where $\Delta k = k_p - k_s - k_i$ is the wavevector mismatch of the involved fields. The generated photons are conventionally referred to as signal and idler in order of decreasing frequency. These conservation laws are the cause of the strong correlations between the generated photons.

Incidentally, the two conservation laws limit the possible wavevector relations between the two emitted photons, known as phase-matching conditions. While these conditions make the emission highly directional, which is advantageous when collecting the generated photons, the limited control one has over the dispersion of the material makes phase matching possible only for specific sets of wavelengths. The most common technique to achieve phase matching is to exploit the birefringence displayed by many crystals. The conditions can be fulfilled in two configurations called type-I phase matching, where the signal and idler field have parallel polarizations but are orthogonal to the pump, or type-II, when signal and idler polarizations are orthogonal to each other. The second-order nonlinerarity is present in many inorganic crystals, such as LN, BBO and KTP, which have large second-order susceptibilities. Conventionally, angle of incidence and temperature tuning are used to eliminate the wavevector mismatch ($\Delta k = 0$). There are, however, scenarios in which the phase-matching conditions for the desired conversion wavelengths may not be achievable by these means. The technique of quasi-phase-matching can be used when conventional phase matching is not possible [103]. This technique involves using a periodically-poled material, a structure that has been engineered such that one of the crystalline axes is periodically inverted as a function of the crystal length. As a consequence, the sign of the nonlinear coupling is also regularly inverted. This introduces an additional wavevector in the phase-matching condition $\Delta k = k_p - k_s - k_i - 2\pi/\Lambda_{pp}$. By choosing the poling period Λ_{pp} such that it matches twice the coherent buildup length of the nonlinear interaction (typically on the order of 10 µm), an inversion of the crystal axis can prevent the field amplitude of the generated wave from decreasing due to a non-zero wavevector mismatch Δk. In this way the field amplitude can grow monotonically. Additionally, quasi-phase-matching allows collinear parametric interaction to be achieved. The inversion of the crystalline axis can be achieved during the fabrication process by applying voltage pulses to the crystal through periodically structured electrodes [104].

In appropriately engineered media, type-0 phase matching is possible as well [105]. Here all three involved polarizations are aligned. This is of special interest for integrating SPDC sources on optical circuit chips and waveguides [106, 78].

We distinguish between a degenerate SPDC process, where signal and idler ideally are identical, and a non-degenerate one, where the two generated photons are

emitted at different wavelengths and can be distinguished. A degenerate down-conversion source requires more degrees of freedom and therefore more effort to be implemented. Degenerate operation is usually chosen when a photon-pair source rather than a heralded source is desired [107].

Photons generated by bulk SPDC typically have a spectral width in the order of hundreds of GHz to THz [108]. Two different approaches can be followed to narrow down the emission spectrum for bandwidth matching an SPDC source to an atomic-ensemble based quantum memory. The first one is to spectrally filter the output of the source. The obvious drawback, however, is that since only a tiny fraction of the spectrum is maintained, the brightness of the source, which is already intrinsically limited to avoid multi-pair emission, is drastically reduced. The alternative is to place the SPDC medium inside an optical cavity, forming an optical parametric oscillator (OPO), which narrows down the emission to the cavity decay rate. Additionally, the cavity can enhance the brightness of the source in comparison to what could be achieved with the same crystal in a non-resonant source design.

Only when the cavity resonances coincide for both wavelengths, can the biphoton propagate in the cavity. In the degenerate scenario this results in a comb like emission spectrum, with the free spectral range (FSR) of the cavity spacing the resonances, modulated by the gain envelope of the nonlinear process. For degenerate SPDC in a cavity, where signal and idler have the same FSR, the spectrum of the biphotons is the product of the downconversion gain envelope and the cavity's Airy transfer function. For the non-degenerate scenario, however, signal and idler experience different FSRs in the cavity, resulting in a so called 'clustering effect'. Each cluster consists of one joint resonance that overlaps for signal and idler and several joint resonances where the modes only partially overlap. After a certain number of FSRs the signal and idler resonances overlap again, forming the next cluster. The spectrum can be computed by multiplying the Airy functions describing the comb-like spectrum for signal and idler by each other and by the SPDC gain envelope. An in-depth theoretical model for the downconversion process in OPOs can be found in [108]. By appropriately engineering the cavity properties it is possible to have one dominating resonance per cluster and in some scenarios even a single cluster within the gain envelope [109]. Except for this last specific scenario, the output of the OPO needs to be spectrally filtered, e.g. with etalons, to have single-mode emission. The separation of adjacent clusters depends on the difference of the group indices for signal and idler and can be approximated by [110]

$$\Delta\nu_{\text{cluster}} = \frac{\Delta\nu_{\text{FSR},s}\Delta\nu_{\text{FSR},i}}{|\Delta\nu_{\text{FSR},s} - \Delta\nu_{\text{FSR},i}|} = \frac{c}{2L}\frac{1}{|n_{g,i} - n_{g,s}|} \,. \tag{2.5}$$

While bulk SPDC can reach very high heralding efficiencies (e.g. 97 % in [107]), when using an OPO this quantity is drastically reduced due to optical losses. In order to still be able to narrow down the emission bandwidth while maximizing heralding efficiency, a monolithic OPO approach can be pursued to reduce losses (see Ch. 3).

The ideal state emitted by a (non-degenerate) downconversion source is the two-mode squeezed state. For small pair generation probability $\lambda \ll 1$ it can be expressed as [111]

$$|\psi\rangle \approx \sqrt{1-\lambda^2}\,|0,0\rangle + \lambda\,|1,1\rangle + \mathcal{O}\left(\lambda^2\right) . \tag{2.6}$$

From the wave function it becomes clear that the emitted photons are exclusively generated in pairs. From the equation above the inherent trade-off of this generation process between emitted-pair rate and multi-pair emission probability becomes obvious. For this reason SPDC sources are operated far below threshold when used for the generation of heralded single photons. This is especially important when they are operated in an OPO, because the presence of the emitted photons could stimulate the process and lead to parametric amplification.

When one of the photons of a pair is discarded or ignored (mathematically speaking, traced out), as is done for example in the measurement of the unconditioned autocorrelations, the remaining mode behaves like the radiation field of a thermal source.

2.2 Quantum Memories

An optical quantum memory can be defined as a stationary device that takes an input quantum state (or part of it, if considering dual-rail approaches) of a single photon and preserves it faithfully in time. Furthermore, this device should be able to re-emit said quantum state at a later point in time on-demand, e.g. when receiving an asynchronous trigger. The qualifier 'optical', which I will omit from here on, refers to the fact that visible and NIR photons are considered. To differentiate this type of device from the read-only memory I presented in the previous section and listed as a SPS, it can be referred to as read-write memory. In general, the quality achieved through the storage should be higher than what could be achieved with purely classical means, e.g. measure-and-prepare schemes.

The general principle of the quantum memory storage process is a coherent mapping of a state of light into a state of matter, and vice versa for the retrieval. A whole plethora of physical systems is currently being investigated for the implementation of quantum memories. The variety of media used for storing single photons (or at least attenuated laser pulses in some examples) includes platforms such as rare-earth doped ions [112, 113, 114, 115], nitrogen-vacancy centers [116, 24, 25], cold atomic ensembles [117, 118, 119, 120, 121], single trapped atoms [122, 123] and, as in our case, atomic vapors at room temperature (or above) [57, 124, 125, 58, 55, 28, 126, 56, 59].Each platform offers particular advantages over others, and it is plausible that there will not be one perfect quantum memory implementation for all applications, but rather, as in the field of classical memories, solutions optimized for specific tasks.

Historically, quantum memory protocols are based on two distinct approaches for storage, which could be categorized as optically-controlled memories and engineered

absorption memories [6]. In optically-controlled memories a strong control pulse is used to mediate storage and retrieval. Since the retrieval is induced by the application of a second control pulse, this type of memory can be read out at arbitrary times. Prominent examples for this approach are EIT, which will be discussed in detail in the following sections, and off-resonant Raman memories.

In the engineered absorption approach the properties of an inhomogeneous broadening are tailored to obtain the desired absorption and re-emission characteristics. This allows for intriguing schemes as for example (longitudinal) controlled reversible inhomogeneous broadening (CRIB) [127, 128] or gradient echo memory (GEM). Here an electric or magnetic field gradient is used to inhomogeneously broaden the absorption line. After absorbing a photon, the atoms start accumulating a phase which depends on the local detuning from the signal frequency. By reversing the external field, the atoms experience a detuning with opposite sign, which leads to a rephasing of the ensemble and successive re-emission of the photon. Another example is the atomic frequency comb (AFC) protocol [115], where the absorption profile is tailored into a periodic comb with equispaced narrow peaks by spectral hole burning. When a photon, which is spectrally broader than multiple comb spacings, is collectively absorbed, the atoms at different frequencies start to dephase. The photon is re-emitted after the rephasing time, which corresponds to the inverse of the comb spacing.

To achieve true on-demand read-out, the engineered absorption protocols can be extended by applying a π-pulse after a photon is absorbed by the medium [115]. The pulse maps the photon to a metastable state, pausing the rephasing of the atomic spin. After an arbitrary delay, a second pulse transfers the atoms back to the excited state, where the precession continues until re-emission occurs. Such protocols can be considered hybrid schemes.

There are two common methods to enhance the light-matter interaction in order to obtain an efficient memory. The first one is to place the medium in an optical cavity, so that the light experiences many round-trips increasing the absorption probability. The second one is to use an optically thick ensemble and exploit the collective coupling [129], which enhances the interaction by a factor $\sqrt{N}$ [130, 131], with N being the number of atoms.

For scenarios were the necessary delay is known in advance, and thus on-demand read-out capability is not necessary, there are easier ways of preserving a photon than mapping it to a matter state. Delay lines can be used to propagate the photon over a predetermined length L, e.g. in free space or through an optical fiber, to achieve the desired delay $\Delta t = nL/c$, with the refractive index of the medium of propagation n. However, considering that for significant delays long delay lines are needed, this technique is affected by loss. A fancy alternative would be to delay the light by significantly changing its group velocity as is done in slow light [132] (see Section 2.2.2). To make them more versatile, delay lines can be looped. For example, a traveling wave resonator with a polarizing beam splitter can be used. By adding a device that can rotate the polarization, e.g. a Pockels cell, the time at

which the photon exits the cavity can be controlled, as an integer multiple of the round-trip time, by rotating its polarization. Empowered with loops, the distinction between delay lines and memories can become blurry. In reference [133] such a delay line was successfully used to synchronize the output of two heralded SPSs. These were operated with a pulsed pump laser whose period τ was matched to a round-trip time of the delay loop. Thus, the first emitted photon could be delayed by an integer multiple of the round-trip time and synchronized with the second heralded photon, even though its emission time was not known *a priori*. The fact that only discrete delay steps are possible, represents no limitation for the task at hand, thus it is appropriate to call the used device a quantum memory in this context.

Analogously to the SPS section, I start by introducing the figures of merit, primarily relying on the definitions provided in Refs. [3, 134]. The figures of merit will again be defined to include all unavoidable imperfections. It is worth mentioning that the quantum memory protocols are not exclusively associated to one platform. Most of them have been demonstrated in more than one physical system, and most platforms can implement more than one scheme. However, so far there has not been a combination that simultaneously excels in all figures of merit. The importance of the figures of merit depends, however, on the specific application, which means that different memories perform better at different tasks.

Fidelity The fidelity is the analog to what the indistinguishability is for a SPS, with the difference that the fidelity is measured between the 'same' photon before and after being stored. For arbitrary states described by a density matrix ρ it is defined as the state overlap between the input state and the one that is read out [2, 135]

$$F = \mathrm{Tr}\sqrt{\sqrt{\rho_{\mathrm{out}}}\rho_{\mathrm{in}}\sqrt{\rho_{\mathrm{out}}}}\,. \tag{2.7}$$

If both input and output states are pure, the expression above reduces to $|\langle\psi_{\mathrm{in}}|\psi_{\mathrm{out}}\rangle|$.

For realistic memories with finite chance of success it is convenient to make this quantity independent from the efficiency with which the memory reads out photons. By conditioning the fidelity on the success of the read-out, where $F = \eta F_{\mathrm{cond}}$, we obtain the conditional fidelity which quantifies the deterioration of state independently of the efficiency [3].

Efficiency The efficiency of a quantum memory is the success probability of storing and re-emitting a single photon [3]. To be more precise, I consider here the end-to-end efficiency, which is the quantity directly accessible by an experiment and is the relevant metric for assessing whether a memory is ready for a certain application. It is the probability to have a photon in the single-mode fiber downstream of the memory after retrieval, if there was a photon as an input. The end-to-end efficiency

is specific for a certain storage time (cf. section memory lifetime). The end-to-end efficiency, also referred to as external efficiency, is defined as

$$\eta_{e2e} = \frac{N_{\text{ret}} - N_{\text{noise}}}{\eta_{\text{source}}\eta_{\text{det}}N_{\text{trig}}} \,. \tag{2.8}$$

Here N_{ret} represents the number of counted photons in the retrieval and N_{noise} the detected read-out noise measured in a separate measurement under the same conditions but with no input. N_{trig} describes the number of attempts to store a photon performed in the experiment, which, depending on the mode of operation of the memory, is given either by the number of received triggers, for a periodically triggered memory, or the number of detected heralds, for a reactive memory interfaced with a probabilistic source. By dividing by the detector and source efficiency this metric is made independent from the imperfections of these two devices. However, other technical losses that are inherent to the successful operation of the memory, for example signal losses in a noise filter, are not factored out [4], since without these inefficiencies the memory would not be able to perform as stated.

For ease of comparison between different memories, the efficiency is often extrapolated to 'short', actually zero, storage time after determining the memory's lifetime. If this extrapolated value is additionally divided by the signal transmission through the setup, the internal efficiency can be obtained from the measured end-to-end efficiency. The internal efficiency describes how efficiently the physical processes of storage and retrieval take place in the considered system. It is the product of the efficiencies for the two separate processes of reading the photons in and out [130], which are the quantities that can be obtained through theoretical models.

Lifetime One possible way of defining the lifetime is through the coherence time of the storage state of the memory [5]. It constitutes the time up to which the collective superposition is preserved before it starts to dephase, e.g. because of coupling to the environment. Dephasing makes the state incoherent, rendering it useless for the intended application.

The lifetime constitutes an upper bound for the maximum storage time after which the memory read-out is still useful for a specific task [87]. This definition is contextual and can be defined as the time after which a specific figure of merit, e.g. the efficiency, has fallen below a certain threshold, commonly $1/e$ [126, 134]. In the following chapters I will report the $1/e$ lifetime of the end-to-end efficiency.

For a fairer comparison of lifetimes across different memory implementations, the fractional delay, defined as the ratio of storage time to pulse width, is an appropriate metric [136]. Since the pulse width is approximately the minimum time a gate operation takes, the fractional delay corresponds roughly to the maximum number of possible operations, e.g. synchronization attempts, that can be performed within the lifetime of the memory [126].

Read-Out Noise An important figure of merit for characterizing the performance of memory implementations prone to read-out noise, as are for example memories in hot vapors, is the signal-to-noise ratio (SNR)

$$\mathrm{SNR} = \frac{N_{\mathrm{ret}} - N_{\mathrm{noise}}}{N_{\mathrm{noise}}} \,. \tag{2.9}$$

N_{ret} represents the number of detected events within a time window in which the retrieval is attempted. N_{noise} on the other hand is an estimation for the read-out noise obtained from a comparable experiment run without signal input. The noise events are counted in the same time window. It can be legitimate to adapt the time gating by narrowing or widening the considered window to better suppress the generated noise, as long as the same region of interest is chosen for all figures of merit for consistent specifications.

The μ_1 parameter is defined as the inverse of the SNR corrected for possible inefficiencies of the signal input [137, 58, 126]. It describes the necessary mean photon number of an input in order to achieve a SNR of one for the characterized memory [94]. This number can be used to set requirements for SPSs to be interfaced with the memory.

Following the definition of the $g^{(2)}$ for SPSs, one can also determine the state accuracy of the output retrieved from the memory. If a single-photon state was stored into the memory, according to the definition of fidelity, a single-photon output is desired. However, the read-out noise strongly affects this quantity by contaminating the retrieval and by altering the measured $g^{(2)}$ value. In fact, measuring the single-photon nature of read-out has been a long-standing problem for broadband ground-state vapor cell memories [54, 55, 56] (more about how this was solved in Ch. 6). To estimate the $g^{(2)}$ of the read-out expected from theory, in [55] the authors have built an incoherent model in which the combined photon statistics of the input photon and the read-out noise are considered. Following the argumentation from [138] they obtain

$$g^{(2)}_{\mathrm{ret,theo}} = \frac{(N_{\mathrm{ret}} - N_{\mathrm{noise}})^2 g^{(2)}_{\mathrm{input}} + 2N_{\mathrm{noise}}(N_{\mathrm{ret}} - N_{\mathrm{noise}}) + N^2_{\mathrm{noise}} g^{(2)}_{\mathrm{noise}}}{N^2_{\mathrm{ret}}} \,. \tag{2.10}$$

This model does not include the effects of all possible noise sources. Especially for coherent noise processes like FWM, a coherent model is necessary [55]. Nevertheless, from the equation above we can get an impression of how the second-order autocorrelation is affected by read-out noise.

Modes In the context of quantum memories, the mode capacity can refer to various concepts. The most intuitive one refers to spatial modes, which actually corresponds to the number of modes that can be stored, for example in an atomic ensemble [139, 140]. In [141] this concept has been dubbed 'quantum hologram'. It can also refer to temporal modes [115, 142]. In this context, a train of identical pulses, only distinguishable by their arrival time, is stored into the memory. Here the mode

capacity can be interpreted as the number of pulses that can be stored before the first one needs to be read out again. Temporal multi-mode memories are of interest for long-distance communication, since they promise higher rates [40].

Compatibility In order to work, a memory needs to be matched to the bandwidth and wavelength of the incoming photons. Additionally, a high bandwidth is intrinsically desirable, since it is the defining quantity for the operation speed in an application [3]. For synchronization applications, the bandwidth is incorporated in the so called time-bandwidth product – the product of the acceptance bandwidth and the memory lifetime. This is the critical metric for synchronization of spontaneous SPSs using quantum memories as it determines the increase in coincidence rates [49]. For long-distance quantum communication, photons in the telecom bands have the great advantage of having minimal absorption losses within optical fibers. However, this does not directly imply that memories must operate at those wavelengths. QFC [143] is making fast developments. This new resource has already proven to be a well suited bridging technology that allows the frequency conversion of single photons preserving, or even enhancing, their indistinguishability [144]. External conversion efficiencies of 57 % have recently been demonstrated [145] for QFC.

2.2.1 Vapor Cell Memories

Alkali metal vapors constitute a good and simple platform for implementing quantum memories [146, 125, 59, 28]. Due to their high number density, high optical depths (ODs) for transitions in the NIR can be reached for temperatures in a range of 50 – 100 °C [6]. Said transitions can be optically pumped and manipulated with diode lasers, making them easily accessible. For such experiments (see [3, 4, 6] and references therein), Rb and Cs are usually chosen, among other properties, for their ground-state splitting of several GHz, which facilitates separation of the signal from the control beam.

The metastable ground states of alkali atoms make for excellent storage states. In fact, long coherence times can be reached. The memory implementations are usually limited by other factors, such as the diffusion of the atoms out of the interaction volume or collisions that perturb the spin state [147, 148, 56].

By cooling and trapping the atoms, the lifetime of the memory can be prolonged. That, however, comes at the expense of an increased experimental complexity, therefore losing one of the key advantages of this platform: the experimental simplicity of vapor cell setups, which should not be neglected when considering the scalability of this system and how realistically it could be operated under real-world conditions.

An important class of memories is based on an ensemble composed of atoms with a possibly isolated three-level system. Two useful configurations of the atomic energy levels are the lambda [147, 146, 149] and the ladder (or cascade) [28, 126] schemes.

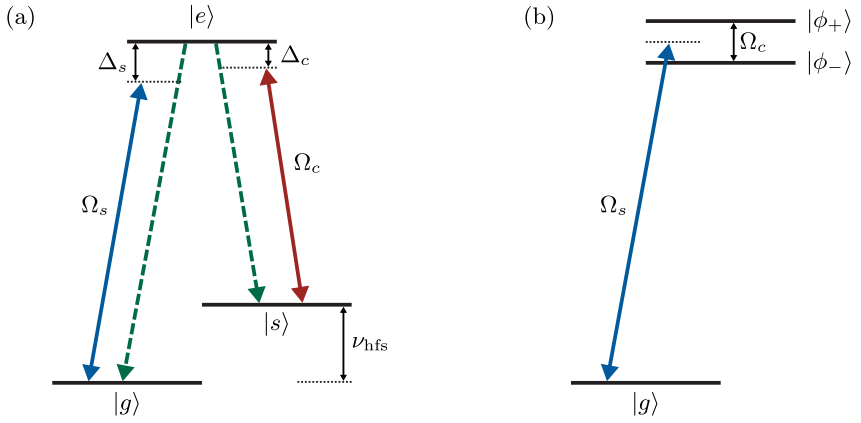

Figure 2.2: Three-level system. Ω denotes the Rabi frequency of the control or the signal, respectively. The latter is assumed to be very small, since we consider a single photon. **(a)** Lambda system as used in EIT based storage protocols. **(b)** Dressed states generated by coherent coupling of the control laser give rise to the AT doublet.

I will proceed by going into detail about the first configuration, as it corresponds to the scheme we implement in our memories, and then briefly comment on the pros and cons of the latter.

2.2.2 Lambda-Memories

One of the most common configurations used for quantum memories is the lambda scheme, named after the Greek letter Λ due to its resemblance to the involved energy configuration (see Fig. 2.2(a)). The level system is composed of a ground state $|g\rangle$ coupled to the excited state $|e\rangle$ through the signal transition with a detuning Δ_s. Preferentially, $|g\rangle$ is chosen to be a dark state for the control, which couples the storage state $|s\rangle$ to $|e\rangle$ with a detuning Δ_c. Usually, the system is kept on two-photon resonance, meaning that the individual detunings are equal ($\Delta = \Delta_s = \Delta_c$), since the considered processes then become more efficient. In alkali vapors $|g\rangle$ and $|s\rangle$ are conveniently chosen to be metastable ground states, whenever possible different hyperfine states, with no optical transition between the two. We define the decay rate as $\Gamma = \Gamma_{eg} + \Gamma_{es}$, where Γ_{ex} describes the decay rate from the excited state into state $|x\rangle$. Initially, the atomic medium is prepared by optically pumping all the atoms into the state $|g\rangle$.

Depending on the detuning Δ, different scattering regimes describe the storage processes taking place. Even though the physical description may vary, in [131] it has been shown that identical maximal efficiencies, which depend only on the OD, can be obtained throughout the different detuning regimes.

Electromagnetically Induced Transparency In the regime where the working point detuning of the memory is comparable to the width of the (broadened)

linewidth, the storage process is best described by the phenomenon of electromagnetically induced transparency (EIT). For better understanding, let us start by first considering the phenomenon for a static control field.

Static EIT causes an opaque medium, which in absence of the control field would just absorb the signal, to become transparent for said signal [150]. For small two-photon detunings $\Delta_s \approx \Delta_c$ the absorption is strongly reduced, while in the case of two-photon resonance full transparency can be reached.

The transparency can be understood as destructive interference between the direct excitation pathway $|g\rangle \rightarrow |e\rangle$ and the indirect pathway $|g\rangle \rightarrow |e\rangle \rightarrow |s\rangle \rightarrow |e\rangle$ (or higher order) in a laser-dressed medium. Due to the control field's much higher intensity, the indirect absorption pathway has a probability amplitude of the same order of magnitude as the direct one, but with opposite sign for resonant fields [151]. The linear electric susceptibility χ determines the refractive index of an atomic medium. Its real part is thus related to the dispersion, while its imaginary term describes the absorption within a medium. The linear susceptibility for the signal transition in the considered three-level system (without inhomogeneous broadening) can be expressed as

$$\chi = -\frac{n_a \sigma_{eg}}{k_{eg}} \frac{2\Gamma_{eg}(\Delta_s - \Delta_c)}{2i\Gamma(\Delta_s - \Delta_c) + 4\Delta_s(\Delta_s - \Delta_c) - |\Omega_c|^2} \,. \tag{2.11}$$

A full description of this phenomenon and the derivation of the expression above can be found in topical reviews as [151, 152]. In the equation above n_a is the atomic number density, σ_{eg} the scattering cross section, and k_{eg} the wavevector of the signal transition. The real and imaginary parts of the susceptibility are visualized in Fig. 2.3(a) as a function of the signal detuning. It is noticeable that the transparency window is accompanied by a (very) strong dispersion, which usually occurs where the absorption is strong. This property can thus be exploited to manipulate the light propagation with the atoms. Recalling that the group velocity of the light follows the relation $v_g \propto (\frac{\partial \chi}{\partial \omega})^{-1}$ [153], EIT can be used to slow down light. Slow light can propagate with significantly reduced group velocity $v_g \ll c$. In [154] the authors reached a group velocity of $17\,\mathrm{m\,s^{-1}}$ in an ultracold atomic ensemble.

The phenomenological description above is valid for the regime where $\Omega_c \ll \Gamma$, usually called the EIT regime. By increasing the Rabi frequency of the control field with respect to the (broadened) linewidth, there is a smooth transition to the Autler-Townes (AT) regime. Note that Eq. (2.11) is valid throughout both regimes. Due to the involved strong control field, the AT regime is best described in the dressed state picture. Neglecting the weak probe field we have a typical dressed two-level system (see Fig. 2.2(b)). On resonance, the dipole interaction Hamiltonian is diagonal in the coupled basis with the eigenstates $|\phi_\pm\rangle = \frac{1}{\sqrt{2}}(|s\rangle \pm |e\rangle)$ with an energy shift of $\pm\frac{\hbar}{2}\Omega_c$ with respect to the energy of the bare excited state. The two peaks in the absorption plot in Fig. 2.3(b) can be understood as the resonances of the two dressed states, known as the AT doublet, induced by the control laser. The energy splitting is $\hbar\Omega_c$.

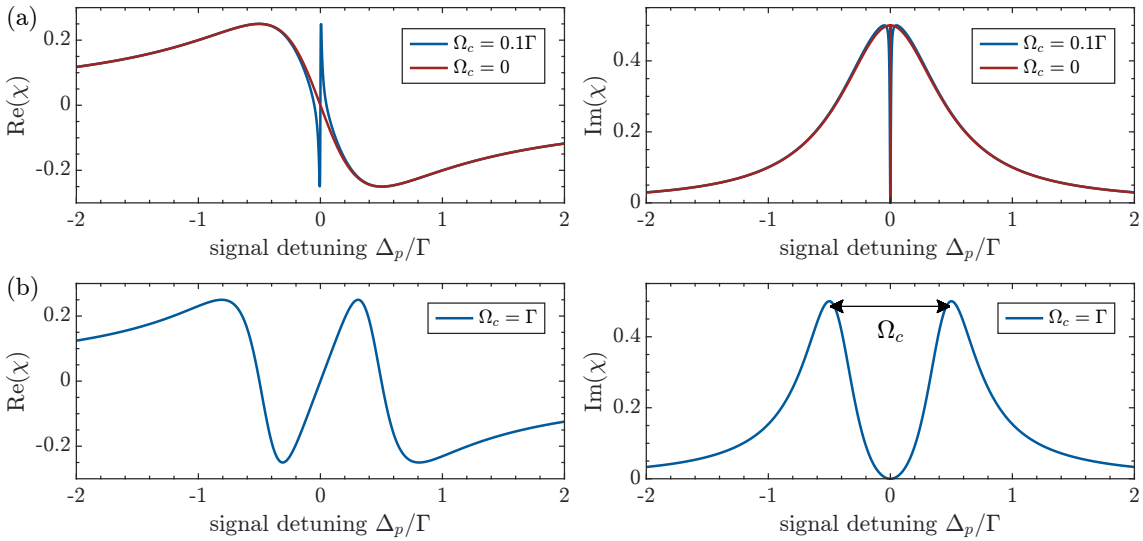

Figure 2.3: Real and imaginary parts of the electrical susceptibility of the signal transition in a three-level system as a function of the detuning. The control field is on resonance $\Delta_c = 0$. The susceptibility is plotted in units of $\frac{n_a \sigma_{eg}}{k_{eg}}$. Two scenarios are considered: **(a)** $\Omega_c \ll \Gamma$ and **(b)** $\Omega_c > \Gamma$, corresponding to the EIT and AT regime respectively. The red curve in (a) shows the susceptibility in absence of the control.

In this regime the following picture can be adopted: if the signal is tuned to the zero-field resonance between the two dressed states, since these coupled states are equally spaced, but with opposite sign of detuning, they have the same contribution of linear susceptibility but opposite sign, which leads to destructive interference of decay channels [151]. In fact, for $\Omega_s \ll \Omega_c$, which is being considered here, $|g\rangle$ becomes a dark state for the signal field. This can be seen as the border case of coherent population trapping, where the two fields are highly unbalanced. The generated dark state, which is a superposition of the two ground states, has an increasing component of state $|g\rangle$ with decreasing Ω_s. A detailed description of EIT in the dark state picture can be found in [152].

By changing the control's Rabi frequency, EIT can be used to store light, as first proposed in [155]. Tuning off the control field ($\Omega_c \to 0$) further reduces the group velocity. If the transparency window is 'closed' while the signal pulse is inside of the atomic medium, the light can be 'stopped'. By turning the control field on once more, the group velocity increases again and the photon exits the ensemble. However, this 'stopping' of the light should not be taken literally, since it is associated with the temporary transfer of the photon's energy to the medium. In fact, the signal is stored as a superposition between the two states $|g\rangle$ and $|s\rangle$, which is often referred to as a spin wave. The system can be described as a combined excitation of the light field and the atomic spin, called a dark-state polariton [156]. The transfer of the quantum state of the signal photon to the atomic ensemble during storage can

be described by

$$|1\rangle_{\text{ph}} \otimes |g_1, g_2, ..., g_N\rangle \longleftrightarrow |0\rangle_{\text{ph}} \otimes \frac{1}{\sqrt{N}} \sum_{j=1}^{N} e^{i\Delta k z_j} |g_1, ..., s_j, ..., g_N\rangle \tag{2.12}$$

where N is the number of atoms, z_j the position of the j-th atom along propagation axis, and Δk the wavevector difference of the control and signal fields [2]. The photonic excitation is mapped onto the atoms as a shared coherent excitation between the states $|g\rangle$ and $|s\rangle$ of the atoms. This Dicke state with a single excitation, which is often referred to as W state, is particularly robust against particle loss [129, 157]. If atoms are lost, the remaining system remains entangled. The spin wave does not contain any contribution of the excited state $|e\rangle$, and it is therefore not affected by its short lifetime. The retrieval is commenced by turning on the control field again, whereupon the collective dipole moment re-emits the stored photon. By applying the read-out control pulse in the reverse direction, one initiates the exact time-reversal of the storage process, called 'backward' retrieval, which yields maximal achievable efficiencies. The photon is thus re-emitted along the same direction it came from, which can pose experimental difficulties. For this reason, and for phase matching considerations as well, it can therefore be more convenient to apply the second control pulse along the same direction, in what is known as 'forward' retrieval. Although this approach usually reaches slightly lower efficiencies, which do not always tend to unity for increasing OD, it can have advantages in terms of the achievable signal transmission through a setup.

The requirement for an optically thick medium can be understood from the necessity to simultaneously fulfill the following two requirements [158]: Firstly, the slowed down pulse needs to fit inside the medium. In order to avoid the leading edge of the pulse to leak through or the tail to be cut off, the pulse length $L_p = v_g t$ in the medium should not exceed the length L of the atomic ensemble. Hence, an upper bound for the pulse duration can be expressed as $t \leq \tau = L/v_g$, with the group delay τ. Secondly, to be efficient, the spectral components of the pulse need to fit inside the transparency window, to avoid absorption and the resulting losses through spontaneous emission. In the EIT regime the FWHM of the transparency window $\Delta\omega_{\text{EIT}}$ in an optically thick medium can be obtained from the susceptibility in the form:

$$\Delta\omega_{\text{EIT}} = \frac{\sqrt{2}|\Omega_c|^2}{4\Gamma\sqrt{\text{OD}}} \,. \tag{2.13}$$

The second requirement can be expressed as lower bound for the pulse duration $t \geq \Delta\omega_{\text{EIT}}^{-1}$. The group delay of the medium can be described in terms of the static EIT window width as $\tau = \sqrt{\text{OD}}/\Delta\omega_{\text{EIT}}$ [158]. The conditions specified above can be combined into the ratios $L/L_p \sim \tau/t \leq \sqrt{\text{OD}}$ [156]. The larger these ratios the better the memory performance. This requires a medium with a high OD $\gg 1$, which corresponds to a significant slowdown of the signal pulse. For optimally matched temporal profiles of signal and control pulses, the OD of the medium is the only parameter that determines the storage efficiency.

Off-Resonant Raman Scattering At this point I shall briefly mention another possible scattering regime in which a memory in the lambda scheme can be operated. For detunings Δ much larger than the broadened linewidth, off-resonant Raman scattering describes the physical process of the light-matter interaction.

A photon with a large detuning Δ is sent into the vapor cell. Simultaneously, a control pulse on two-photon resonance is applied. The photon is absorbed via a virtual state induced by the control, and a Stokes photon is emitted. The photonic excitation is mapped, analogously to the EIT scenario described earlier, to a shared coherent excitation between the two involved ground states. For the retrieval process, a second control pulse with the same detuning is applied. The chosen detuning can be significantly larger than the bandwidth of the photon, so that the excited state is never occupied.

The dressed state picture can be used to understand the Raman memory scheme as well. In this scenario, the driving field is detuned from the bare resonance. The coupled states are thus asymmetrically split, one close and one further away from the resonance. The incoming signal is not tuned between the two dressed states, as in EIT, but onto resonance with the state $|\phi_-\rangle$, which has the larger $|s\rangle$ component [159]. The fundamental difference to the EIT scheme is that there is no interference of absorption pathways in this protocol, and that there is no group velocity reduction of the light in the atomic ensemble [3].

Bandwidth, Noise, and Lifetime As we have seen, quantum memories in warm atomic ensembles allow for high-bandwidth storage, which is not, as one might think, limited to the rather narrow natural linewidth of the atomic species. The physical limit is given by the splitting between the two states $|g\rangle$ and $|s\rangle$, commonly chosen to be of two different manifolds of the hyperfine split ground state. When the bandwidth of the photon to be stored becomes of the same order of magnitude as this splitting, it couples to both states and the storage process starts failing. The practical limit, however, usually is constituted by the available control power. In fact, the acceptance bandwidth of the memory depends, as we have seen, on the Rabi frequency Ω_c of the control transition, which is defined by $\hbar\Omega_c = \mathbf{dE}$, where $\mathbf{d}$ is the dipole moment of the transition (e.g. for the ^{87}Rb cycling transition, which is the strongest, $d = 2.53 \times 10^{-29}\,\mathrm{C\,m}$) and $\mathbf{E}$ the electric field of the control laser. For a beam of approximately 0.5 mm in diameter, an optical power of at least 100 mW is necessary for the control field to drive the cycling transition on the D_2 line (which has the highest transition strength) with a Rabi frequency of $\sim 1\,\mathrm{GHz}$. From an experimental point of view, this sets the requirement of using optical amplifiers when working with diode lasers. Experiments involving successful storage of attenuated laser pulses with bandwidths exceeding 1 GHz have already been reported [149].

The necessary intense control pulses expose one of the main challenges of implementing a quantum memory: separating the single signal photon from the control. Various filtering techniques (cf. Ch. 5) such as polarization, spectral, and in some cases spatial filtering, help suppress the control, while attenuating the signal as little

as possible. However, while the control can successfully be suppressed by 16 orders of magnitude (see Ch. 6) the read-out pulse can also induce noise processes in the atoms, which result in the emission of spurious photons that cannot be filtered out from the signal, since they have the same frequency and polarization. Commonly, three main noise-generating processes are identified: fluorescence noise, stimulated Raman scattering, and FWM noise [6].

Fluorescence Noise – Let us assume for the discussion of the noise mechanisms that the storage state $|s\rangle$ is chosen to be in the other hyperfine ground-state manifold with respect to $|g\rangle$. The two states are thus split by the hyperfine splitting $h\nu_{\mathrm{hfs}}$. For warm Rb $\nu_{\mathrm{hfs}} \approx 10\Gamma$, where Γ is the Doppler-broadened linewidth. Even if off-resonant, the control field can couple to the ground state and induce fluorescence. Due to collisional relaxation in the excited state[1], this fluorescence occurs at the natural frequency of the spectral line with a Doppler-broadened profile, even when induced by a detuned field [161]. The resulting light is thus dubbed 'collisional' fluorescence. By working with a detuning of the signal, but still on two-photon resonance, this type of fluorescence can be filtered out. Due to the Doppler-broadening, however, large detunings are necessary, leading to reduced coupling strengths. Due to their far off-resonant working point, this type of noise is less of an issue for memories in the Raman scheme. An additional way to mitigate the effect of this noise channel is time gating. The D_2 excited state in Rb has a lifetime of approximately 26 ns. If the stored signal pulses are significantly shorter than this lifetime, then the contribution of the fluorescence can be notably reduced. In [57] this method has been demonstrated in a Cs memory.

Stimulated Raman Scattering – The thermal population distribution of the ensemble is given by the Boltzmann factor, and unless $h\nu_{\mathrm{hfs}} \gg k_B T$, all the ground states are equally populated. For alkali vapors with hyperfine splittings $< 10\,\mathrm{GHz}$, this is true at room temperature. If the atoms are not actively removed, the control pulses will scatter population in $|s\rangle$, creating anti-Stokes noise (still considering the atomic configuration from Fig. 2.2(a)). These photons have the same frequency and polarization as the signal. This process is known as stimulated Raman scattering. Since the ground-state splitting is, to some degree (see Ch. 4), dictated by the atomic species and the high temperatures are necessary to reach high ODs in the vapor, the ensemble needs to be optically pumped to a good degree to empty the state $|s\rangle$ before a storage attempt is undertaken.

FWM Noise – The last noise process is FWM, which is schematically shown in Fig. 2.4. The control field off-resonantly couples to the ground state $|g\rangle$, producing a Stokes photon and transferring an atom to $|s\rangle$. A second photon from the control pulse can then excite this atom, which emits an anti-Stokes photon. While the Stokes photon can be filtered, the anti-Stokes photon has the same properties as the signal, and cannot be filtered. Often the statement is found, that for $\nu_{\mathrm{hfs}} \gtrsim \Delta$ the memory interaction couples more strongly than the spontaneous FWM and

[1] When these collisions happen during the write or read processes they can also lead to decoherence [160].

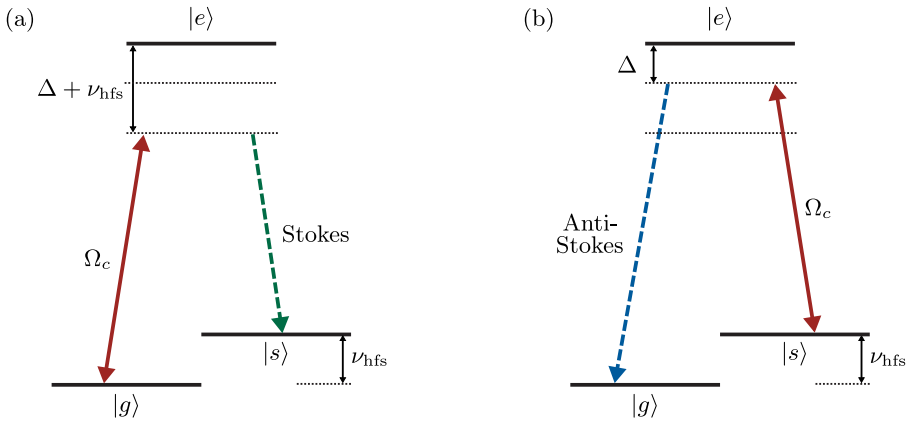

Figure 2.4: Four-wave mixing. **(a)** The control field couples off-resonantly to the ground state $|g\rangle$, generating a Stokes photon transferring an atom to $|s\rangle$. **(b)** The control field can now couple $|s\rangle$, producing an anti-Stokes photon that cannot be filtered since it has the same properties as the signal.

that only for far off-resonant Raman protocols does FWM become a significant source of noise [6]. However, this statement might be misleading. In [162, 163] the authors found that FWM can dominate even on resonance if the OD is high enough. Unfortunately, a high OD is desirable for an efficient memory. The model for the field propagation inside an EIT medium with FWM developed by [164] determined the functional limit for an EIT quantum memory. According to their findings, the noise becomes exponentially large for OD $> \Delta|\Omega'|/\gamma_{eg}|\Omega|$, where $|\Omega|$ is the control field Rabi frequency, $|\Omega'|$ is the Rabi frequency of the detuned control field on the signal transition, and γ_{eg} the decoherence rate. In order to be able to reach high enough ODs to make the storage process efficient, the ratio γ_{eg}/Δ needs thus to be minimized. Another reason to limit FWM is that it lowers the fidelity of the memory.

The noise considerations introduced so far assume a three-level system. Real atoms, however, have many more states that need to been taken care of in an application. Additional states can lead to spurious processes that lower the memory performance. For example, in scenarios where the signal couples to an excited state for which the control transition is forbidden, the photon is just absorbed and lost. Or, if there is more than one excited state coupling to the ground states, the coupling might interfere and cancel out. Selection rules, optical pumping and external fields are helpful tools to isolate (near) ideal systems that I will explore in the chapters to come.

Limitations of Memory Lifetime – To conclude the discussion about lambda-scheme memories in hot vapors, I will consider dephasing and decoherence effects that affect the memory's lifetime. EIT itself has a strong angular dependence in Doppler-broadened atomic vapor [165]: for storage, an angle between the beams can lead to dephasing. In Eq. (2.12) we see that the wavevector difference Δk between signal and control beam causes a spatial phase dependence in the spin wave. If the atoms move during the storage time, which is the case for a warm ensemble, the spin wave

dephases and the memory performance decreases upon retrieval. The wavelength of the spin wave $\lambda_{SW} = 2\pi/\Delta k$ should be longer than the vapor cell's length, so that the phase mapped onto all atoms is homogeneous [166]. The maximal spin-wave wavelength is obtained for a collinear configuration. The dependence on the angle between the two beams is strong enough, that angles large enough (compared to the typical length scale of an optical setup) for even mild spatial filtering become the limiting factor of the lifetime in hot vapor memory implementations. In [56] an angle of only 10(1) mrad was sufficient to limit the lifetime to 68(7) ns. In cold atom implementations, where the atoms move significantly slower, this poses a minor problem, as in [167] for example the authors used an angle of 13° and reached a lifetime on the order of 4 µs. The expression for determining the spin-wave dephasing-limited memory lifetime can be found in [166, 167]. A Gaussian decay of the memory efficiency is expected for dephasing [166].

In hot vapors and in the absence of a buffer gas, the atoms fly through the vapor cell performing a ballistic motion with an average velocity $\langle v_T \rangle = \sqrt{8k_BT/\pi m}$. They thus fly in and out of the interaction volume interrogated by the field. The transit time of an atom through a beam with radius r can be estimated by $\tau = 2r/\langle v_T \rangle$. A time τ after the write pulse, the atoms with the shared excitation have moved out of the interaction volume and are no longer addressed by the read-out pulse. Alternatively, if the vapor cell diameter is matched to the transverse beam profile, so that the atoms are confined within the interaction volume, at most after a time τ the atoms collide with the cell walls, randomizing their spin state. To mitigate the effect of wall collisions, anti-relaxation coatings can be applied to the vapor cell to preserve the spin states (see Section 4.4.2). In order to reduce the fast thermal ballistic motion of the atoms, often an inert buffer gas with a small collisional dephasing cross-section for alkali ground states is added to the vapor cell. This changes the atoms' motion to a diffusive one, resulting in a longer atom-laser interaction time given by [158]

$$\tau^{-1} = 2.405^2 \frac{D}{r^2 + 6.8\lambda_{\text{mfp}} r} \tag{2.14}$$

where D is the diffusion constant, and $\lambda_{\text{mfp}} = 3D/\langle v_T \rangle_{\text{alk-bg}}$ the mean free path between collisions of the alkali atoms with the mean relative velocity $\langle v_T \rangle_{\text{alk-bg}} = \sqrt{8k_BT/\pi M}$. Here M is the reduced mass of the colliding system.

Another relevant decoherence mechanism originates from alkali-alkali collisions, which can lead to pair-wise spin exchange. This type of collisions involves the swapping of the electron spins of the collision partners, conserving the total angular momentum, which leads to scrambling of the Zeeman state of the involved atoms. The rate of decoherence through spin-exchange γ_{SE} is given by $\gamma_{SE} = n\sigma_{SE}\langle v_T \rangle_{\text{alk-alk}}$, with the alkali number density n, the spin-exchange cross-section σ_{SE}, which for Rb is $\sigma_{SE} = 1.9 \times 10^{-14}\,\text{cm}^2$, and the relative mean velocity between the two alkali atoms $\langle v_T \rangle_{\text{alk-alk}}$ [168]. For the so called 'stretched' state, where the spin is fully aligned, spin exchange is forbidden due to conservation of angular momentum. By preparing the atomic ensemble in this particular states this decoherence channel can be suppressed. Another way to suppress spin-exchange induced decoherence is to

work in a particular environment of low static magnetic field and high alkali density known as spin-exchange relaxation free (SERF) regime. This regime is often used for the operation of vapor cell magnetometers [169] and was successfully applied in a quantum memory implementation, which using anti-relaxation coatings as well, demonstrated a lifetime of 430 ms in hot Cs vapor [59]. Furthermore, it has been observed that in high magnetic fields (kG) the Rb-Rb relaxation rates are reduced by a factor of 3 even at multiatmosphere buffer gas pressures compared to the zero-field scenario [170].

2.2.3 Ladder-Memories

Another interesting three-level scheme for the realization of quantum memories is the cascaded or ladder scheme. Here the storage state is chosen to be an excited state of a higher electron orbital instead of a ground state. This bears the obvious drawback of the spin-wave coherence time being fundamentally limited by the excited state's radiative decay lifetime. Additionally, due to the significantly different wavelengths of the two transitions, the spin wave dephases after an even shorter time because of the wavevector mismatch. However, the scheme comes with attractive advantages as well. The difference in frequency between the two addressed frequencies is so large that it becomes trivial to suppress the control field with simple color filters. The physical constraint for the bandwidth is, technically speaking, lifted. In the scenario with $|g\rangle$ and $|s\rangle$ separated by the hyperfine ground-state splitting, the photons are limited to bandwidth smaller than this splitting, so that they cannot accidentally couple to both ground states. For ladder-memories this does not constitute the liming factor due to the very large splitting, allowing for almost arbitrarily high bandwidths. Furthermore, only very little noise is generated since there is no population in the storage state chosen to be doubly-excited. No pumping is thus required, meaning that no preparation time is necessary before each storage attempt, which in our implementations poses the current limit to the repetition rate of the experiment.

Ladder-memories have been implemented both in Cs [28] and in Rb [126, 171]. In both cases the storage state was chosen to be in the D orbital. In this scheme no FWM can occur and noise from imperfect pumping is not an issue. Apparently, collisional fluorescence seems not to be a problem either for this scheme, since in both demonstrations no narrow-band filtering was necessary. In [28] the authors were able to successfully store and retrieve heralded single photons with no added noise.

In order to protect the collective state from inhomogeneous dephasing, techniques like dressing the storage state with an auxiliary state have been developed [172]. Despite the short storage times, these memories are envisioned to work as 'noise-free quantum buffer' to make identical photons from distinguishable noisy SPSs by filtering their temporal-spectral mode [173]. Recently, a ladder-memory implemented in Rb has been demonstrated to be able to directly store photons with a

wavelength in the telecom C-band [174]. Despite the very limited fractional-delay, this implementation marks an interesting development.

2.3 Numerical Simulation

In hot vapors it is challenging to maintain high transparency and long ground-state coherence for high number densities [158]. It is therefore practical to optimize storage for an imperfect but real scenario with moderate ODs by seeking a balance between losing the tails of the pulse and having a finite spectral EIT width. Gorshkov *et al.* [131, 175, 130, 176, 177] have published a series of theoretical works on lambda-scheme storage, valid for EIT, off-resonant Raman and some photon-echo memory protocols, both for the scenario of atoms in a cavity [175] and in free-space [130], including a treatment for inhomogeneous broadening [176]. This work represents a useful tool for optimizing the performance of a quantum memory and even offers insights on how to apply a gradient-ascent algorithm to perform the optimization numerically [177]. In this series of works the authors found that there is a maximum achievable efficiency that exclusively depends on the OD, since the photonic coupling to the spin wave depends on the OD. Furthermore, achieving maximal efficiency does not require an adiabatic turning on or off of the control pulse and can be achieved by optimizing the control shape to the signal input or vice versa. The optimization procedures described in the aforementioned theoretical work have been successfully demonstrated for the case of moderate OD < 25 [148]. For increasing ODs, on the other hand, a deviation from theory was observed, attributed to FWM being a competing process and having an impact on the spin wave, as they described in [54].

As a tool to further improve our quantum memories after a proof-of-principle demonstration, I implemented a numerical simulation of the storage and retrieval process based on the aforementioned theoretical works. The simulation is meant to guide us through the vast interdependent parameter space for analyzing limitations of the experiment and guiding future development.

For the simulation I considered an atomic four-level system following the description in [178], which itself builds on the free space treatment of Gorshkov *et al.* [130, 131]. A uniform distribution of atoms within the interaction volume is assumed. In our experiments, the condition posed in [130] that the control beam be much wider than the transverse mode of the signal is not always met. For this reason, I decided to take into account the transverse profile of both light fields by implementing the three-dimensional equations of motion, as used in [179] and derived in detail in [159]. I considered the scenario of collinearly propagating signal and control beam in the paraxial approximation. This allowed me to simplify the considered problem by exploiting the cylindrical symmetry through integration over the angle ϕ. In order to simulate a hot vapor I included both collisional broadening induced by the buffer gas as well as inhomogeneous Doppler broadening. To model the latter, different velocity classes of the atoms are introduced, as shown in [176]. The atoms are sorted

into j separate velocity classes, such that all atoms in a certain velocity class have approximately the same velocity v_j an thus experience the same detuning. It is assumed that the individual atoms do not change their velocities during the storage and retrieval processes. This assumption is valid as long as these processes are much faster than the mean free time between collisions. In the vapor cell used for the Zeeman memory, the mean free time between collisions is approximately 20 ns, thereby fulfilling this requirement. However, during the considered storage times of > 80 ns it is assumed that the atoms fully rethermalize, which means that the atoms fully redistribute among the different velocity classes. By increasing the number of velocity classes the computation becomes more accurate, albeit at the cost of new equations of motion, making the simulation more time-consuming. While the effect of the Doppler-broadening is considered in the simulation, the actual motion of the atoms is neglected – the atoms are considered stationary but experiencing a detuning. The motion of the atoms plays a crucial role for the dephasing of the spin wave. However, for simplicity, no deterioration of the spin wave during the storage time is considered. Both backward as forward retrieval can be implemented in the simulation. Because of its relevance to the experimental approach, I focused on the latter scenario.

In the current implementation, the simulation does not include any noise model. The results must thus undergo a critical sanity check. A typical pitfall would be that the simulation finds the highest efficiency for the memory on resonance, which is also where the read-out noise would be the highest – this would most certainly not constitute the best operation point.

2.3.1 Equations of Motion

In this section I will present the equations of motion implemented for the simulation and describe how the efficiencies of the storage and retrieval process are calculated. I will, however, refrain from reproducing the derivation of said equations and refer the reader to the corresponding sources. The equations of storage and retrieval in a four-level system are derived in [178], the implementation of velocity classes for treating inhomogeneous broadening in [176] and the derivation of the three-dimensional equations of motion for a three-level system can be found in [159].

A lambda scheme is considered as shown in Fig. 2.2(a). However, to approximate a more realistic scenario, a second excited level that can couple to both the signal and control fields is added. The two excited states have a frequency splitting Δ_e. The incoming photon is described by a quantum field with a slowly varying envelope operator $\mathcal{E}(z, \rho, t)$. In order to describe the involved dynamics, three further sets of operators are necessary: the polarization operators $P_{1,j}(z, \rho, t)$ and $P_{2,j}(z, \rho, t)$ for the respective excited state and the spin wave operators $S_j(z, \rho, t)$. The index j denotes the velocity class. The intense control field, with a peak amplitude $\Omega_{\max}$, is implemented as a classical field $\Omega(z, \rho, t)$, which is considered not to be affected by the storage and retrieval interactions. Both the transverse profile as well as the

temporal profile of the control pulse are programmed to be a Gaussian distribution. Both its waist and temporal width are freely selectable. Both light fields are assumed to be centered at the corresponding transition frequency with a detuning Δ_s and Δ_c respectively.

The equations of motion are derived under the dipole and rotating wave approximations. The quantum field is considered to be weak. This implies the assumption that almost all atoms remain in $|g\rangle$ during the process and only terms to linear order in $\mathcal{E}$ are retained. The incoming quantum field is assumed to have only one non-empty mode in the time interval $[0, T]$. This incoming mode is mapped onto some mode of S during the storage process. The retrieval is defined as the reverse process of mapping back the spin wave mode to an outgoing field at a later point in time $T_r > T$.

For modeling the inhomogeneous broadening, the Gaussian velocity distribution (along the optical axis) of the atoms for the considered ensemble temperature is divided into classes with equal velocity span, for a total of j velocity classes (preferably an odd number to include the zero-velocity class). The fraction of atoms in each class is expressed by p_j, with $\sum_j p_j = 1$. The central velocity of each class v_j is used to compute the associated Doppler shift $\Delta_j^{\mathrm{D}} = \omega_0 v_j / c$, where ω_0 is the unperturbed transition frequency.

Expressed in a coordinate system where $\tau = t - z/c$ is the time in a co-moving reference frame and $\tilde{z}$ is a dimensionless length ($\tilde{z} \in [0, 1]$), the equations of motion are

$$\left(\frac{\alpha^2}{2ik_s}\partial_\rho^2 + \partial_{\tilde{z}}\right)\mathcal{E} = i\sqrt{d\gamma}\sum_j \left(\sqrt{p_j}P_{1,j}\mu_{1,g} + \sqrt{p_j}P_{2,j}\mu_{2,g}\right) \tag{2.15}$$

$$\partial_t P_{1,j} = -(\gamma + \gamma_{\mathrm{hom}} + i(\Delta_s + \Delta_j^{\mathrm{D}}))P_{1,j} + i\sqrt{d\gamma p_j}\mathcal{E}\mu_{1,g} + i\Omega S_j \mu_{1,s} \tag{2.16}$$

$$\partial_t P_{2,j} = -(\gamma + \gamma_{\mathrm{hom}} + i(\Delta_s + \Delta_j^{\mathrm{D}}) - i\Delta_e)P_{2,j} + i\sqrt{d\gamma p_j}\mathcal{E}\mu_{2,g} + i\Omega S_j \mu_{2,s} \tag{2.17}$$

$$\partial_t S_j = i\Omega^*(\mu_{1,s}P_{1,j} + \mu_{2,s}P_{2,j}) + i(\Delta_s - \Delta_c)S_j \ . \tag{2.18}$$

The transverse derivative in Eq. (2.15) describes the diffraction of the signal field as it propagates. The aspect ratio $\alpha = L/w_s$ is introduced to correctly scale the two spatial coordinates. Here, the length of the vapor cell is denoted by L and w_s represents the signal beam waist. The magnitude of the signal wavevector is given by k_s. γ is the decay of the coherence of the excited state, corresponding to half of the spontaneous emission rate $\Gamma = 2\gamma$ (for a purely radiative decay). Homogeneous broadening, e.g. as the collisional broadening due to buffer gas, is accounted for by γ_{hom}. We adopted the notational conventions from [178]. As is common in theoretical works, the resonant OD and the Rabi frequency Ω_{max} are defined as half of their experimental value. To avoid confusion, I introduce the quantity $d = \mathrm{OD}/2$. The relative dipole moments $\mu_{x,y}$, where $x = 1, 2$ denotes the considered excited state and $y = g, s$ the ground state, are expressed in units of the dipole moment of the

two-level cycling transition. Consequently, $\Omega_{\max}$ and d are expressed for the cycling transition as well. This (apparently) facilitates the comparison between different systems.

For the computation, dimensionless time and frequency quantities in units of the FWHM of the control pulse or its inverse value are used respectively. At the beginning of the storage process, right before it enters the ensemble, the photon is described by the temporal envelope $\mathcal{E}_{in}(\rho, \tau)$. The transverse profile is programmed to be a Gaussian distribution. Furthermore, we assumed a perfectly pumped scenario at the beginning, meaning that all the atoms are initially in the ground state $|g\rangle$, thus there are no excitations in the spin wave nor in the atomic polarizations. The initial boundary conditions for the storage process can be thus expressed as

$$\begin{aligned} \mathcal{E}(\tilde{z} = 0, \rho, \tau) &= \mathcal{E}_{in}(\rho, \tau) \\ P_{i,j}(\tilde{z}, \rho, \tau = 0) = S_j(\tilde{z}, \rho, \tau = 0) &= 0 \qquad \text{with } i = 1, 2\,. \end{aligned} \tag{2.19}$$

The boundary condition for the transverse profile is that for $|\rho| \to \infty$ all involved fields and polarizations vanish. This is achieved by manually programming $\mathcal{E}$ to be zero at the maximum value of ρ considered in the numerics. The system of partial differential Eqs. (2.15)–(2.18) is solved numerically. In the next section I will briefly comment on how I have implemented the numerical computation. Based on these results, the storage efficiency can be computed by taking the ratio of excitations mapped to the spin wave to the incoming photon

$$\eta_s = \frac{2\pi \int_0^\infty \int_0^1 \rho |S_{out}|^2 \,\mathrm{d}\tilde{z}\,\mathrm{d}\rho}{2\pi \int_0^T \int_0^\infty \rho |\mathcal{E}_{in}|^2 \,\mathrm{d}\rho\,\mathrm{d}\tau} \tag{2.20}$$

where $S_{out} = S(\tilde{z}, \rho, \tau = T)$ is the spin wave mapped to the atoms at the end of the storage process. The spin wave S to which all S_j average after rethermalization is given by $S(\tilde{z}, \rho, \tau) = \sum_j \sqrt{p_j} S_j(\tilde{z}, \rho, \tau)$.

The scenario of forward retrieval can be described with the same system of equations with different initial conditions. We consider that at the retrieval time $T_r > T$ the atoms have rethermalized and the spin wave is the same across all velocity classes. At the beginning of the retrieval process there is no excitation in the quantum field $\mathcal{E}$ and the polarizations are also assumed to be zero. As the initial value for the spin wave we take the spin wave resulting from the storage simulation at time T wheigted by the fraction of atoms in the corresponding velocity class. The boundary conditions for the retrieval (denoted by the superscript r) thus are

$$\begin{aligned} \mathcal{E}^r(\tilde{z} = 0, \rho, \tau) &= 0 \\ P^r_{i,j}(\tilde{z}, \rho, \tau = T_r) &= 0 \qquad \text{with } i = 1, 2 \\ S^r_j(\tilde{z}, \rho, \tau = T_r) &= \sqrt{p_j} S_{out}\,. \end{aligned} \tag{2.21}$$

By solving the system of differential equations Eqs. (2.15)–(2.18) with these new initial conditions we can compute the retrieved field $\mathcal{E}^r_{out}(\rho, \tau) = \mathcal{E}^r(\tilde{z} = 1, \rho, \tau)$.

The retrieval efficiency is obtained analogously to what we have seen for the storage process

$$\eta_r = \frac{2\pi \int_{T_r}^{\infty} \int_0^{\infty} \rho |\mathcal{E}_{out}^r|^2 \, \mathrm{d}\rho \, \mathrm{d}\tau}{2\pi \int_0^{\infty} \int_0^1 \rho |S_{in}^r|^2 \, \mathrm{d}\tilde{z} \, \mathrm{d}\rho} \, , \tag{2.22}$$

where $S_{in}^r = S_{out}$.

For comparisons with experimental data, where it is difficult to assess the efficiency of the two separate processes, it is convenient to compute the internal efficiency, which in theoretical publications is referred to as total efficiency, and is given by

$$\eta_{tot} = \eta_s \times \eta_r \, . \tag{2.23}$$

Alternatively, one could directly compute the ratio between the output and input $\mathcal{E}$ fields.

2.3.2 Some Comments on the Numerical Methods

Our problem expresses as a system of $3j + 1$ partial differential equations as a function of time and two spatial variables, as shown in Eqs. (2.15)–(2.18). I started implementing the simulation in Matlab having some of the original code of [178] at my disposal. Considering the increased amount of differential equations to solve compared to the stationary problem, due to the treatment of the inhomogeneous broadening, it became soon obvious that the performance was a critical point of the simulation that needed improvement. A simple approach to solve the differential equations with the Euler method would not suffice anymore – some more advanced numerical techniques would be required to compute the problem within a reasonable amount of time and an acceptable numerical precision. This fact motivated the search for better suited numerical algorithms, which finally converged on the following technique: to solve the time-dependent equations I recurred to a Runge-Kutta method and used the method of spectral collocation to solve the equations depending on the spatial variables. In the following, I will summarize the key points of the numerical methods I implemented to efficiently solve the problem at hand and suggest to the interested reader to have a look at the appendix of [159] for a pedagogical description of these numerical algorithms.

For finding numerical solutions, the time and space variables need to be discretized. The numerical simulation is based upon the method of lines. The solution is computed by iteratively stepping forward in time, starting at $t = 0$. The spatial derivatives are computed at once for all positions, while the temporal derivatives are computed incrementally. The method of lines allows for a coarse discretization of the spatial coordinates. The solution of the computation appears to "trace out" a line at each spatial point, which describes the temporal solution for that position in space, hence the name of the method.

For the computation of the spatial derivatives I use the technique of spectral collocation using Chebyshev polynomials. This technique can be explained most intuitively through a Fourier transform. For periodic problems, the Fourier transform can be used to convert differentiation into ordinary multiplication. The linearity of the Fourier transform allows us to transform, perform the multiplication, and to take the inverse transform, expressing the full differential equation as a matrix multiplication. This concept can be generalized to non-periodic problems. For solving such problems, non-periodic basis functions – polynomials here – are used instead of complex exponentials to approximate the function of interest. This can be understood as follows: A vector f (representing a function sampled at the collocation points) is fitted with a polynomial p. The derivative of the polynomial p' can be easily computed. Evaluating p' at the collocation points gives the vector approximating the derivative of the function f. This property can be used to construct the differentiation matrix D [180].

Using a polynomial differentiation matrix with Chebyshev points (the projections of equispaced points along the arc of a semicircle) as collocation points is a convenient choice for most non-periodic problems [181]. Chebyshev points cluster towards the boundaries of the domain. This fact brings an advantage by avoiding the instability of equally spaced grid points, which can lead to oscillations near the edges of the domain, known as Runge phenomenon [181]. For smooth functions, this method is very accurate, even if only a few spatial points are used. By approximating the continuous function $A(z)$ as vector elements sampled at the grid points $A_j = A(z_j)$ (and B analogously) this technique allows us to simplify the problem to an algebraic equation that can be solved in one run:

$$\partial_z A(z) = B(z) \qquad \longrightarrow \qquad D\mathbf{A} = \mathbf{B} \tag{2.24}$$

where D is the Chebyshev differentiation matrix. This equation can easily be solved for A, either by multiplying the inverse of D, or by using the 'backlash operator' (relying on Gaussian elimination) when working with Matlab, which is significantly faster. Implementing the boundary conditions can become a little bit counterintuitive with spectral collocation, since they need to be directly incorporated in the dynamical equation by modifying the differentiation matrix and the first entries in B. To compute the Chebyshev grids and the corresponding differentiation matrices for this simulation I used the function `cheb.m` described in [180].

The temporal derivatives appearing in the equations of motion, on the other hand, are solved with a partially-implicit second-order Runge-Kutta method. For understanding how these two techniques are combined and how to expand the problem to two spatial and one temporal dimensions, I point the reader to [159].

Relying on a recursive calculation, the simulation *per se* cannot be parallelized. However, the algorithm is implemented so that when an input parameter is being swept, the computation of the independent values of that parameter is parallelized on different computing cores of the computer. The size of the grid used for the computation has major effects on both the computational time required as the accuracy

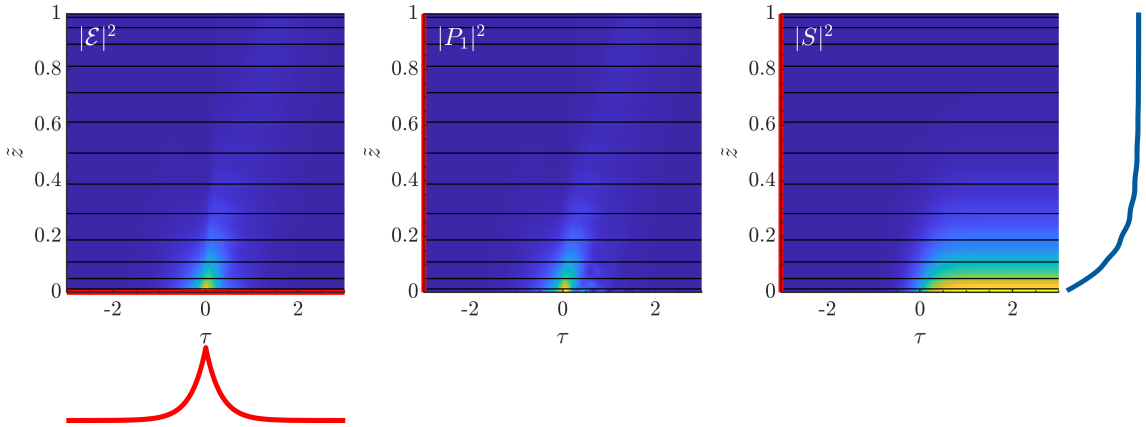

Figure 2.5: Sample simulation for the storage process. The intensities of the computed results are shown for the center of the transverse profile. The x-axes are given in units of the temporal FWHM of the control pulse. The boundary conditions are shown in red. The temporal profile of the incoming photon $\mathcal{E}_{in}$ is shown underneath the plot. The blue curve on the right shows the spin wave stored across the atoms at the end of the storage process. This distribution represents the input for a retrieval simulation.

of the results. To estimate a good trade-off, I compared the results of the simulation for an increasing number of grid points. I found that a typical grid size to reach convergence is composed of 1500 points along the temporal axis, while 35 points along each spatial direction suffice. The primary limiting factor for the number of implemented velocity classes, on the other hand, is dictated by the available RAM of the used computer.

2.3.3 Example Solution for the Simulation

To give a feeling for the results obtained by the simulation described above, in Fig. 2.5 I show the output fields for the storage process. In order keep it simple, the visualized scenario is computed with all the atoms in one velocity class, which corresponds to the stationary-atoms case with no detuning due to the Doppler effect. The plots show how the fields of interest evolve along the center of the transverse profile of the beams ($\rho = 0$). The red lines indicate the boundary conditions implemented in order to obtain these solutions. The initial condition for $\mathcal{E}$ corresponds to the temporal envelope of the photon to be stored. To keep the example scenario simple, the couplings of the fields to the second excited state were set to zero. The blue line next to the panel visualizing S shows the spatial distribution of the spin wave once the storage process is concluded. All the excitation left in the signal field $\mathcal{E}$ at the end of the ensemble ($\tilde{z} = 1$) is not mapped to the spin wave and can be experimentally recognized as the signal leakage through the cell. The horizontal lines in the plots represent the collocation points chosen for the computation of the spatial coordinate. It can be noted that towards the edges of the space domain the points lie closer together, following the distribution of the Chebyshev points. In

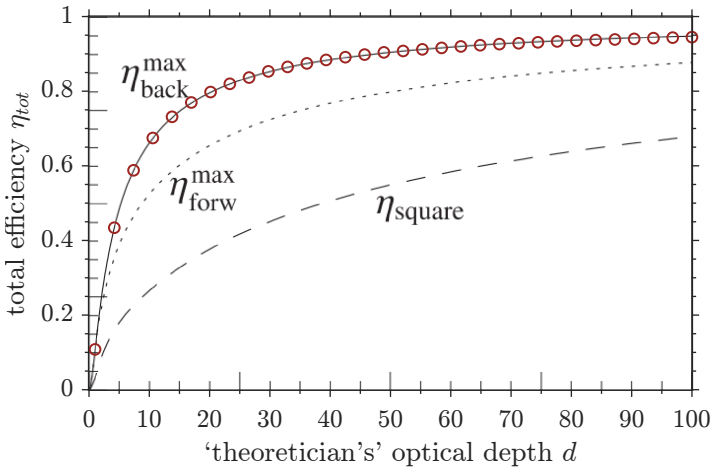

Figure 2.6: Maximal achievable efficiencies for optimized pulse shapes as a function of the 'theoretician's' optical depth d. The red circles show the outcome of my simulation, run iteratively with the gradient ascent algorithm from [178] to optimize the control's temporal profile. The data points are overlaid to Fig. 4 reproduced from [130]. Reprinted figure with permission from A. V. Gorshkov *et al.*, PRA **76**, 033805 (2007). Copyright 2024 by the American Physical Society.

order to obtain the smooth plots, a cubic `spline` interpolant is used along the z direction.

As a sanity check of the results provided by the simulation I wrote, I cross-checked them against the maximum total efficiency derived in [130]. For this purpose I adapted the gradient-ascent algorithm from [178] to be compatible with my simulation, and computed the efficiency for an optimized control pulse for different ODs. Figure 2.6 shows the efficiency obtained from my computation for optimized pulses as red circles overlaid with Fig. 4 from [130]. The agreement is very good, and constitutes a validation of the implemented code. Because of some issues I ran into regarding the convergence of the optimization algorithm, I performed these comparison for the scenario of backward retrieval.

3 Heralded Single-Photon Source

A long-term goal of our in-house collaboration with the Warburton group is to interface one of their Rb-like quantum dots with an atomic vapor memory to build a hybrid quantum node. In 2017, after [56] had been published, it was clear that both systems still needed to be improved in order to reach this goal.

As an intermediate step we decided to characterize the performance of the quantum memories we are working on with 'real' single photons from another, more readily available, source. A tailored SPDC source, designed to be compatible with hot vapor memories, appeared to be a promising test-bed. For pursuing this goal we teamed up with Oliver Benson's group at the Humboldt-Universität zu Berlin who shared their expertise on these systems with us. The original design and assembly of the source was led by Janik Wolters, who joined Benson's group in Berlin for this purpose. Around mid-September 2018 the first version of the source was shipped over $\sim$870 km from Berlin to Basel. Mounted on an optical breadboard with a size of $90 \times 60\,\mathrm{cm}^2$ (experiment control computer and laser controller were shipped separately), it could be considered to be 'fully portable'. Once the source arrived in our laboratory I took over its operation, eventually rebuilding it due to the non-linear crystal at its core breaking, more than once. Together with Gianni Buser, who was of great help with the collection and processing of the statistical results, we carefully characterized the photon-pair source's output and its statistics, which we reported in [64].

After losing several crystals, GB and I started discussing how to improve the source with the focus on the ease of operation and longevity of the crystals. I made the necessary adaptations in the design. Shortly after our newly coated batch of crystals arrived, Björn Cotting joined our team. In the frame of his Master's thesis work, he implemented this second iteration of the source based on our design, under the supervision of GB and myself. The full results of his thesis are freely available [182]. Finally, with only some minor adaptations to BC's setup, we were able to successfully interface the source with the Zeeman-scheme memory [29]. These adaptations were mainly needed for suppressing crosstalk between the two systems.

The photon-pair source was specifically designed to operate on the ^{87}Rb D_1 line. In an attempt to push the boundaries of the tuning range of the improved source design we found a working point where the phase-matching conditions are also fulfilled at the D_2 line. This allowed us to use the same source for testing the performance of the hyperfine Paschen-Back memory with single photons as well.

The rest of the chapter starts with the description of the original, triple-resonant source design and then proceeds with the double-resonant implementation, focusing on the changes needed for interfacing with our quantum memories.

3.1 Triple-Resonant OPO

Our goal was to tailor a single-photon source working on this principle to an atomic vapor memory operating at the Rb D_1 line. To achieve this goal the source needs, among other things, to be bandwidth- and wavelength-matched to the memory it is interfaced with. In the following I will elaborate on the aspects that were considered in the original source design.

3.1.1 Design

Quantum memories in hot atomic-vapor ground states can be compatible with single photons with a bandwidth of hundreds of MHz to GHz [149]. Bulk SPDC, however, has an emission spectrum that is in the order of hundreds of GHz to THz (or a few nm) broad [183]. To match the bandwidth of the single-photon source with the acceptance bandwidth of the atomic memory, the non-linear crystal can be placed into an optical parametric oscillator (OPO) and pumped far below threshold. As we have seen in Ch. 2, the linewidth of the emitted photons then depends on the cavity decay rate. Furthermore, in cavity-enhanced SPDC, the brightness of the source is increased. A joint resonance in the spectrum occurs only where energy conservation is satisfied and both signal and idler fields are resonant. If the FSR of the two fields differs, as in our case, the emitted spectrum is composed of clusters of joint resonances instead of a full comb.

To achieve wavelength matching, the source should emit signal photons at the ^{87}Rb D_1 line. Furthermore, a tunability of the emission of at least 1 GHz is necessary in order to have enough flexibility on the memory's side for finding a good working point. For a thorough investigation of red and blue detuned working points of the memory, even larger tunability would be desirable. Given the clustered emission spectrum, it is important that only the heralds corresponding to the desired signal frequency are detected. For this purpose the idler photons need to be spectrally filtered. The signal photons consequently inherit the idler's mode purity for every conditioned measurement. Uncorrelated noise in the signal arm is expected, constituted of the signal photons whose herald has been filtered out or lost. Most of these photons do not disturb the memory operation since they are off-resonant with respect to the targeted atomic transition and thus do not interact with the vapor. Furthermore, these off-resonant signal photons are filtered downstream by the read-out filtration stage before they can reach the detectors and they therefore do not contribute to the measured noise. What does, however, appear as uncorrelated noise in the data are resonant signal photons whose idler was lost. The frequency stability

of the downconverted signal affects the two-photon detuning in the lambda-scheme used for the memory and determines whether a source-memory frequency lock is required. For a source stability comparable to that of a free running laser, the time scales involved in the storage and retrieval process are short enough that an operation without frequency lock should be possible.

A monolithic OPO approach is attractive for a number of reasons. The lack of intra-cavity interfaces reduces losses, while a cavity made out of one piece is intrinsically robust against vibrations, temperature and pressure fluctuations. Double-resonant monolithic [184, 185, 186, 187] and triple-resonant semi-monolithic OPOs [188, 189, 190] have been studied for the generation of squeezed light states. Investigations about the stability and tuning by varying the temperature and the pump frequency have been been performed for double-resonant monolithic waveguide resonators in [191, 192]. The achieved heralding efficiencies and pair generation rates, however, do not reach those required for a successful pairing with a quantum memory in hot atomic vapor.

In order to preserve the non-classical characteristics of the single photon even after retrieval, the SNR of the memory and the heralding efficiency of the source should be compatible. A high heralding efficiency is crucial, since even small amounts of noise can accumulate in the retrieval time window if during most storage attempts no photon is present [94]. Typical SPDC sources with external cavities reach limited heralding efficiencies being limited by the losses on the interfaces within the cavity. By choosing a monolithic design, where highly reflective coatings are directly applied on the crystal surfaces, we reduce the number of optical interfaces to a minimum. Since the alignment of the two mirrors in a monolithic cavity is highly dependent on the manufacturing precision and can't be corrected afterwards, we chose a hemispherical cavity design. This is insensitive to tilting of the curved mirror [193]. The plane surface is chosen to be the output facet of the cavity with the aim to minimize the wavefront distortion of the emitted signal photons, for achieving a better fiber coupling efficiency.

To establish triple resonance in the cavity, three degrees of freedom are necessary. The advantages of a monolithic cavity design come at the cost of losing some standard degrees of freedom: We cannot displace a cavity mirror to vary the resonator length, nor can we tilt the crystal inside the OPO. As degrees of freedom to fulfill the resonance conditions we can tune 1) the periodically-poled potassium titanyl phosphate (ppKTP) temperature, which acts both on the resonator length as well as on the refractive index of the medium, and 2) the pump laser frequency, since we are aiming at a non-degenerate downconversion process. For efficient downconversion the latter should be varied in accordance to the pump resonances within the crystal. The only restriction we have on the pump wavelength is that the resulting idler wavelength must lie within the sensitivity range of our standard detectors and various laboratory equipment. As a novel, albeit unidirectional, way of tuning the cavity we apply strain on the non-linear crystal with a piezoelectric actuator deforming the refractive index ellipsoid, which, with some effort, allows us to reach the cavity resonance conditions.

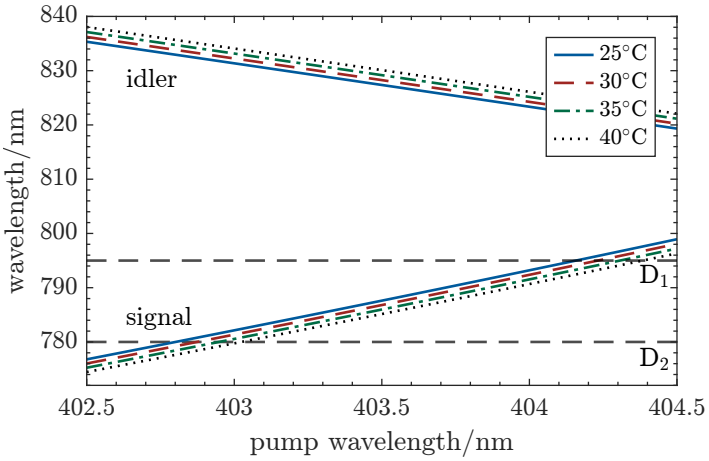

Figure 3.1: Quasi-phase-matching solutions enforcing energy and momentum conservation as a function of the pump wavelength for different crystal temperatures and a poling period $\Lambda_{pp} = 10.1465\,\mu\text{m}$, which better matches the experimental data (see Sec. 3.4). With the chosen poling period the crystal has an operating temperature range of 30 – 50 °C. The temperature-dependent Sellmeier equations as published by S. Emanueli and A. Arie in [194] are used for this calculation. The two dashed horizontal lines represent our target frequencies for the ^{87}Rb D_1 and D_2 line, respectively. The solutions are for bulk SPDC; no cavity is considered at this stage.

Conveniently, the whole fabrication process of the crystal could be outsourced to commercial manufacturers. *Raicol Crystals* manufactured a 28 mm long ppKTP crystal for non-degenerate type-II SPDC for us. The crystal has a cross-section of $1 \times 2\,\text{mm}^2$. The poling period Λ_{pp} of the KTP crystal defines where the quasi-phase-matching conditions can be met. For a given target signal wavelength and a range of working temperatures it can be computed by solving the linear equation system consisting of energy and momentum conservation. Figure 3.1 shows the solutions, signal and idler wavelength pairs, for different crystal temperatures and the nominal poling period of 10.1 µm (with a small correction, see caption), which we chose for fabrication. One can see that for various pump wavelength and crystal temperature configurations, signal photons at 795 nm should be obtained. The crystal was subsequently shipped to *Photon LaserOptik* to be cut into pieces of different lengths (3.0(1) mm, 5.0(1) mm, 7.0(1) mm) and polished to our specifications. Finally, the crystals were dielectrically coated by *Lens-Optics*. The detailed specifications of the polishing and the coatings can be found in Tab. 3.1.

Let us now have a look at how these specifications reflect on the cavity properties of the source. In order to determine the expected bandwidth, we first compute the theoretical finesse, given by the well known relation

$$\mathcal{F} = \frac{\pi}{2\arcsin\left(\frac{1-\sqrt{\rho}}{2\sqrt[4]{\rho}}\right)} \tag{3.1}$$

where ρ is the fraction of circulating power left after a cavity round-trip, which can

Symbol	Property	Surface	Specified value
	Surface quality	plane, curved	10/5 scratch/dig
		top, bottom	20/10 scratch/dig
	Parallelism (deviation)	plane, top, bottom	$<20''$
r	Radius of curvature	curved	10.0(1) mm
	Centricity (deviation)	curved	$<3'$
R_p^p	Reflectivity at 404 nm	plane	99.0(5) %
R_p^c		curved	87.0(5) %
$R_{s/i}^p$	Reflectivity at 795 nm & 820 nm	plane	91.5(5) %
$R_{s/i}^c$		curved	>99.9 %

Table 3.1: Table summarizing the surface and coating properties of the monolithic triple-resonant OPO as specified by the respective manufacturers.

be expressed in terms of the mirror reflectivities and the single-pass losses within the cavity medium $\rho = R_1 R_2 T_{sp}^2$. Before polishing and coating, an absorption coefficient $\alpha(\lambda = 405\,\text{nm}) = 1\,\text{dB}\,\text{cm}^{-1}$ was measured in Berlin for our crystal. The single pass transmission is expected to behave like $T_{sp} = 10^{-\alpha L/10}$, for α in $\text{dB}\,\text{m}^{-1}$.

The absorption of ppKTP at near-infrared wavelengths is low. Small absorption coefficients with negligible effects on the cavity properties have been reported in [195] for a crystal grown by the same manufacturer. We did not perform such a measurement when we first received the crystal. After the dielectric coating is applied on both facets, the effort of extracting the internal losses of the crystal from a measurement would be disproportionately large.

Let us now consider the 7 mm long crystal that was mainly characterized in [64]. From Eq. (3.1) we obtain the theoretical finesse for the signal/idler cavity of $\mathcal{F}_{s/i} = 70(4)$ for a single-pass transmission equal to unity. For the pump cavity we have a transmission $T_{sp} = 85.1(2)\,\%$ resulting in a finesse $\mathcal{F}_p = 13.3(3)$.

The FSR is defined by $\Delta\nu_{\text{FSR}} = c/(2n_g L)$ where n_g is the group index $n_g = n - \frac{\partial n}{\partial \lambda}$. The refractive index n of the ppKTP crystal is temperature and wavelength dependent and can be obtained by the empirical Sellmeier equations from [194]. For a crystal temperature of 30 °C the group indices for the ordinary and extraordinary polarization (corresponding to signal and idler in our case) are $n_g^s \approx 1.91$ and $n_g^i \approx 1.80$, respectively. This results in a $\Delta\nu_{\text{FSR,s}} = 11.2(2)$ GHz and $\Delta\nu_{\text{FSR,i}} = 11.9(2)$ GHz. It is important to be careful with the nomenclature in order to avoid confusion. From the relation $\mathcal{F} = \Delta\nu_{\text{FSR}}/\delta\nu$ we can extract cavity damping rates for the signal and idler of $\delta\nu_s = 160(9)$ MHz and $\delta\nu_i = 170(10)$ MHz respectively. The bandwidth of the emitted photons (or "biphoton bandwidth"), however, will be narrower. In fact, since the resonance conditions for signal and idler need to be fulfilled simultaneously in our triple-resonant OPO, the effective linewidth will correspond to the product of two Lorentzian lines with the respective cavity damping rates. The resulting width

L (mm)	$\mathcal{F}_{s,i}$	$\Delta\nu_{\mathrm{FSR,s}}$ (GHz)	$\Delta\nu_{\mathrm{FSR,i}}$ (GHz)	$\Delta\nu_{\mathrm{cluster}}$ (GHz)	$\delta\nu_s$ (MHz)	$\delta\nu_i$ (MHz)	$\delta\nu_{\mathrm{DRO}}$ (MHz)	T_{sp} (%)	$\mathcal{F}_p$
7.0(1)	70(4)	11.2(2)	11.9(2)	193(3)	160(9)	170(10)	106(6)	85.1(2)	13.3(3)
5.0(1)	70(4)	15.7(3)	16.6(3)	271(5)	224(13)	238(14)	148(9)	89.1(2)	16.5(4)
3.0(1)	70(4)	26.1(9)	27.7(9)	450(15)	373(24)	396(25)	247(16)	93.3(2)	21.8(7)

Table 3.2: Theoretical expectations of the cavity parameters for the triple-resonant OPO for different crystal lengths. The derivation of these values is described for the 7 mm scenario in the text.

$\delta\nu_{\mathrm{DRO}}$ can be obtained by the closed-form expression [196]

$$\delta\nu_{\mathrm{DRO}} = \sqrt{\frac{1}{2}}\sqrt{\sqrt{\delta\nu_i^4 + 6\delta\nu_i^2\delta\nu_s^2 + \delta\nu_s^4} - \delta\nu_i^2 - \delta\nu_s^2}\,. \tag{3.2}$$

From the specified design properties we would expect a $\delta\nu_{\mathrm{DRO}} = 106(6)$ MHz, where DRO stands for double-resonant OPO. Using Eq. (2.5) a cluster separation of about 193 GHz (corresponding to approximately 390 pm) can be computed from the group index difference for the two polarizations. This separation should limit the number of clusters within the downconversion gain envelope to three. Table 3.2 contains the theoretical predictions of the cavity properties for all three crystal lengths we used as triple-resonant OPO.

The heralding efficiency is defined as the probability of having a signal photon (in the optical fiber after the source) conditioned on the detection of an idler. Consequently, losses in the herald arm do not affect the heralding efficiency, as they merely reduce the rate at which we can run our experiments. However, the efficiency is highly dependent on signal losses by the optical elements and the chosen reflectivities for the cavity mirrors. We designed our cavity with asymmetric mirror reflectivities in order to achieve a preferred forward (away from the pump) emission direction. The probability of a signal photon leaving the cavity through the plane facet is given by the extraction efficiency [110]

$$p_{s,\mathrm{for}} = \frac{1 - R_s^p}{1 - R_s^c R_s^p T_{sp}^2}\,. \tag{3.3}$$

Assuming no intra-cavity losses for near-infrared wavelengths, we estimate the probability that a photon is lost through backwards emission to be $p_{s,\mathrm{back}} = 1 - p_{s,\mathrm{for}} = 1.07\,\%$ for the specified reflectivities.

3.1.2 Experimental Setup

Figure 3.2(a) shows the experimental setup of the triple-resonant OPO used as heralded single-photon source. A grating-stabilized external cavity diode laser (ECDL, Toptica DL pro HP) at 404 nm is used to pump the non-linear crystal. It is operated

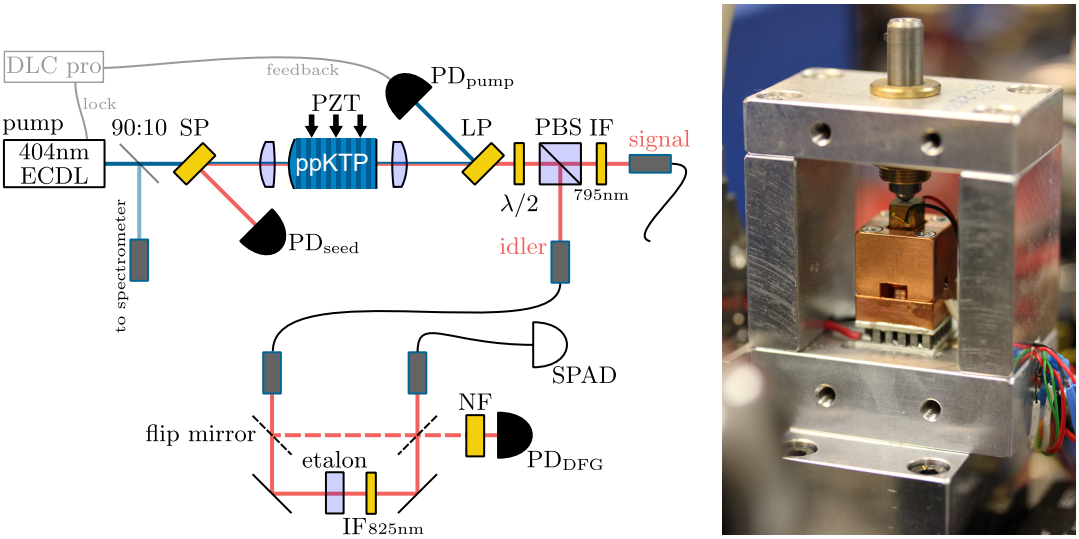

Figure 3.2: Setup of the triple-resonant source. **(a)** Experimental setup of the monolithic triple-resonant OPO. The photon pair is generated by downconversion of a 404 nm photon from the pump laser. Signal and idler photons have orthogonal polarizations in type-II SPDC, which we exploit to separate them with a polarizing beam splitter. While the signal is directly sent to a detection stage, usually an HBT setup, the idler is filtered down to one single spectral mode before detection. For alignment and tuning purposes the OPO is seeded by a tunable laser and the crystal cavity resonances are measured on PD$_{seed}$. For this purpose the seed laser is coupled to the setup in counter-propagating direction through the signal fiber. Used abbreviations: 90:10 - beam sampler with the specified branching ratio; SP (LP) – optical short-(long-)pass filter; PZT – piezoelectric actuator; PBS – polarizing beam splitter; IF – interference filter; NF – 785 nm notch filter; PD – photodiode; SPAD – single-photon avalanche diode. **(b)** Photograph of the oven (housing partially removed) enclosing the crystal. The ppKTP OPO is situated inside the copper mount in the center of the picture. Mechanical strain can be applied by a piezoelectric actuator which can be pre-strained with a fine-threaded screw. The crystal temperature is stabilized by a TEC.

by a digital laser controller (Toptica DLC pro), which allows a larger mode-hop-free tuning range by clever tuning of multiple parameters at once.

The ppKTP crystal is sandwiched between two copper mirrors. These highly polished plates allow for a homogeneous pressure distribution, enabling strain to be safely applied to the crystal without the risk of cracking it. A piezoelectric actuator (Thorlabs PK2FMP1) is used to apply strain from the top, orthogonally to the optical axis. A screw running through the top of the oven mount applies a pre-strain to the actuator in order to allow for a linear tuning range. The voltage for the piezoelectric actuator is provided by the analog output of a computer-controlled 'multifunction I/O device' (NI USB-6002) which is amplified twenty-fold by a high-voltage amplifier (Falco Systems WMA-005). The crystal temperature is controlled by a thermoelectric cooler (TEC or Peltier), driven by a commercial TEC-controller (Meerstetter Engineering TEC-1091). The temperature typically can be stabilized to fluctuations $< 5\,\mathrm{mK}$ at around $30\,^{\circ}\mathrm{C}$. All components are enclosed in an aluminum oven, acting also as heatsink for the TEC. Two small apertures grant optical access. Figure 3.2(b) shows a photograph of the assembly. The crystal mount is positioned between two achromatic doublets. Their focal lengths are chosen to mode-match the incoming pump beam and the outgoing signal and idler modes to their respective modes within the cavity.

For alignment and tuning purposes of the heralded single-photon source a distributed feedback (DFB) laser (TEM Lasy 795) at the target wavelength, the seed, is coupled to the OPO in counter-propagating direction through the signal fiber. This laser beam permits us to align the optical path of the single photons to be emitted and to find the working parameters for tuning the emission to the target wavelength. The seed transmission through the crystal cavity is picked up as the reflection on the short-pass filter and monitored on $\mathrm{PD}_{\mathrm{seed}}$. A CCD camera (Thorlabs, DCC1645M) is used to identify the fundamental transverse mode in the transmission spectrum. Higher modes are suppressed by alignment and good mode-matching of the incoming free space beam with the cavity mode. Furthermore, together with the pump field, the seed laser is used to stimulate the parametric process, resulting in difference frequency generation (DFG). This process generates photons in the same mode as the idler but at a higher emission rate – bright enough to be detected with an amplified photodiode (PD). For aligning the idler arm, a quarter-wave plate is placed between the polarizing beam splitter (PBS) and the ppKTP crystal, so that the seed light backreflected from the plane cavity mirror experiences a polarization rotation of 90° and is reflected by the PBS. This allows the pre-alignment of the idler's fiber coupling. However, since signal and idler wavelengths are about 30 nm apart, it is necessary to optimize the focus of the idler fiber coupler at a later stage. This optimization can be done once a DFG signal is established.

The pump laser is coupled into the OPO through the curved surface of the crystal. Pump light leaking through the plane surface of the cavity is filtered with a spectral long-pass filter (AHF F47-750) to reduce background noise. This filter is mounted on a flip mount as an incoupling aid. Coupling in the pump through the curved facet turns out to be somewhat tedious due to the diverging backreflection. Overlapping

the transmitted pump beam with the previously aligned seed beam significantly simplifies this procedure. Once the alignment is complete, the long-pass filter, flipped back in place, is also used as a pick-up for the transmitted pump light. This allows us to measure the pump cavity transmission on an amplified photodiode with switchable gain (Thorlabs PDA100A2). This signal is used as a feedback to stabilize the laser's frequency with a dither/lock-in amplifier lock and a PID-controller both integrated in the laser controller. In order to distinguish the different pump cavity resonances, and to make the tuning procedure more efficient, a small fraction of the beam is picked up right after the laser and directed to a spectrometer (moglabs economical wavemeter MWM).

An 810 nm half wave-plate is used to rotate the polarization of the photon pair before splitting it, with the intention to send the signal photons through the PBS in transmission, in order to exploit the better polarization extinction ratio. A PBS is used to separate the orthogonally polarized signal and idler photons that are then each coupled into an anti-reflection coated polarization-maintaining (PM) fiber. Due to the cluster structure of the emitted spectrum, we filter the idler photon before we detect it. We use a narrow-band, 8 mm thick, monolithic etalon (FWHM = 274(4) MHz) to suppress all modes within the main cluster except for the central one, and an 825 nm interference filter (IF, Laseroptik) with 0.57(5) nm FWHM to suppress all other clusters. Spectral filters as the abovementioned etalon also play a crucial role in the filtration of the memory read-out and are described in detail in Ch. 5. The idler filtering stage reaches a peak transmission of $\approx 80\,\%$ while its suppression is the product of the extinction of its single elements. The IF reaches an extinction $< 10^{-2}$ for frequencies at least one linewidth away from the resonance, and the etalon reaches a suppression of 1.2×10^{-3} for the FSR/2 frequencies. A more detailed description about how we characterized the filtering stage follows in Section 3.1.6. Since the storage and retrieval measurements we perform are conditioned on the detection of an idler, the idler's spectral purity is transferred to the signal photons. Therefore there is no need to filter the signal photons as stringently. For the measurements performed with the triple-resonant OPO, however, we added a 795 nm IF in the signal arm. Finally, the photons are detected with single-photon avalanche diodes (SPADs, Excelitas SPCM-AQRH-16).

For the alignment procedure the seed beam acts as a proxy for the signal photons and is assumed to have the same wavelength as well as spatial mode. The signal photons are going to be generated in the same cavity mode as the seed laser is coupled to and analogously the idler photon will overlap with the DFG mode. While this is technically true, we experienced that for the single-photon operation the fiber couplings needed slight readjustment to reach peak performance.

Tuning and Measuring the Output Frequency In order to tune the signal output of the source to the desired frequency, the crystal temperature is tuned so that one of the seed's cavity transmission resonances is close to the target frequency. To be able to compensate for possible shifts during the tuning procedure it is conve-

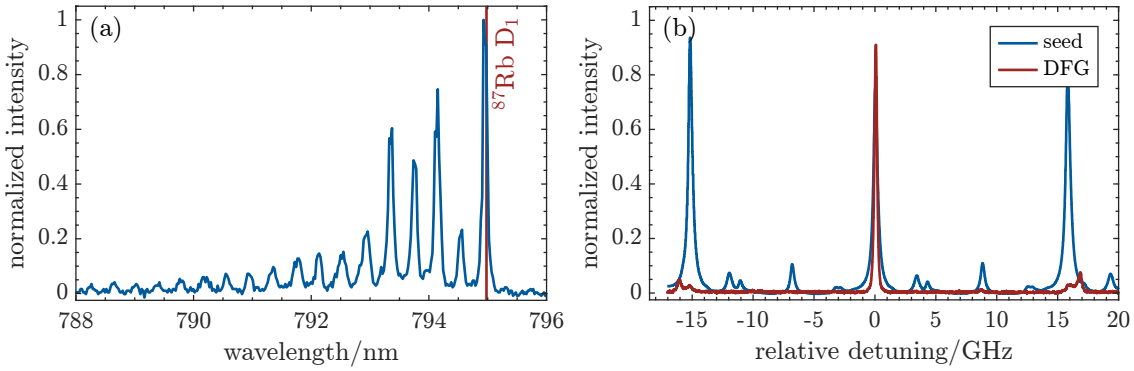

Figure 3.3: Spectral properties of the triple-resonant OPO with a 5 mm crystal. **(a)** Coarse spectrum of the signal photon clusters measured with a 500 mm spectrometer with a CCD camera sensitive to single photons. The spectrum is calibrated with a second measurement where an attenuated Rb D_1 laser at 794.979 nm (vertical line) is sent to the spectrometer. **(b)** Transmission spectrum of the ppKTP cavity (blue), recorded with PD_{seed}. The seed laser is scanned around the ^{87}Rb D_1 line. The three major peaks correspond to the cavity's TEM_{00} modes. The smaller features that are visible are not fully suppressed higher-order modes. The FSR is 15.5(2) GHz and matches the expected value. The corresponding DFG emission, measured simultaneously on the amplified PD_{DFG}, is shown in red. The fine structure of one emission cluster can be seen: a central peak and two doublets separated by an FSR. Outside of the displayed frequency range no other DFG features were measured. The main contribution of the emission lies in the central peak.

nient to start with the seed resonance slightly blue-detuned, since the piezoelectric actuator acts in a unidirectional way.

Subsequently, the pump laser is tuned in frequency until DFG is induced in the crystal. The DFG signal is monitored on an amplified, switchable-gain PD_{DFG} (Thorlabs PDA36A) that can be reached through a flip mirror. A 785 nm notch filter before the PD_{DFG} blocks any residual seed light. Once the pump frequency that generates the DFG signal at the right frequency is found, the pump laser can be frequency locked to its closest resonance. Some iterative fine tuning with the piezo and minor temperature adjustments, on the order of tens to hundreds of mK steps, help reach the target frequency. Finally, the spectral filter stage is tuned to maximally transmit the resulting idler wavelength, thereby cleaning its output to a single mode. A second flip mirror after the spectral filters allows for an easy estimation of the transmission through the filter stage. By unplugging the seed laser at this point, the source emits downconverted single-photon pairs with the signal at the target frequency.

When operating the source for the first time, it is convenient to check whether the emission of the desired wavelength lies within the brightest emission cluster (and if not to correct it by temperature tuning). For this purpose, the spectrum of the SPDC signal emission has been recorded with a 500 mm spectrometer with a CCD camera sensitive to single photons. Figure 3.3(a) shows the clusters of the signal

spectrum of a 5 mm crystal. This measurement was performed with a pump power of about 1 mW. The frequency axis of this measurement has been calibrated with a second measurement performed with an attenuated laser tuned to the Rb D_1 line at 794.979 nm. The spectrometer does not have the resolution to resolve the the spectral structure within each spectral cluster. These data were recorded in Berlin. The presented measurement was not repeated for the other crystals used in the triple-resonant design, since they originate from the same growth process and it seems likely that the operating parameters be the same. We did, however, record an equivalent spectrum for the second iteration of the source (see Ref. [182]), where the same asymmetry was present. The spectrum in Fig. 3.3(a) does not express the typical $sinc^2$-shaped gain envelope of the spectral distribution of the SPDC process (recall Ch. 2), in fact we do observe an asymmetry. In his thesis, Andreas Ahlrichs [108] developed an expression for the joint spectral amplitude of a monolithic, linear triple-resonant cavity. In such an optical resonator the downconversion process can take place in two opposite directions. With the parametric interaction being a coherent process, the relative phase, over which we have no control in our cavity design, leads to interference, which affects the spectrum of the emitted single-photon pairs. The quantum state of the pairs leaving the cavity in forward direction is the superposition of two processes: 1) the photons are generated in forward direction and exit the cavity after a certain number of round-trips, or 2) the photons are generated in the backward direction and after a given number of round-trips leave the cavity. In the latter scenario the pump acquires an additional phase factor, compared to the first scenario, due to the reflection on the flat mirror, and the pair due to reflection on the curved mirror. If necessary, control over this relative phase is usually gained by using either a compensation crystal or specific dielectric coatings with appropriate phase shifts. However, having control over the relative phase seems to be mainly a concern when working in the degenerate scenario. In Fig. 1.17 of [108] a monolithic design with non-poled regions acting as compensation crystals is proposed.

For resolving the spectral structure within one cluster, the most convenient way is to use the signals measured on PD_{seed} and PD_{DFG} during DFG operation as a proxy. Figure 3.3(b) shows a cavity spectrum of the brightest cluster of a 5 mm crystal for the seed laser (blue) and the corresponding DFG signal (red) as a function of the seed frequency. Some higher-order modes on the seed cavity spectrum can be seen. Through a better alignment and mode-matching we typically managed to achieve cleaner spectra than the one in Fig. 3.3(b). Nevertheless, I chose to show this measurement, since it was taken with the same crystal as the spectrometer measurement shown in panel (a).

An interesting technique to characterize the full structure of the SPDC source's emission at once is to use a cascaded filter system as described in [108]. It can be set up by a stack of monolithic etalons that are simultaneously temperature tuned. These etalons have proven to be reliable and easy to operate and have thus found many applications in the various experiments in our laboratory. By choosing resonators with FSRs that are not a multiple integer of each other, the filter system

can have a clean transfer with only one dominating resonance in the frequency range of interest. Furthermore, this resonance basically inherits the linewidth of the narrower filter. This 'optical spectrum analyzer' can thus offer the necessary frequency resolution to resolve the sub-cluster structure as well as a scan range of hundreds of GHz for scanning over various clusters. In order to calibrate the temperature ramp necessary for tuning the system, the FSR of the single etalon needs to be determined precisely, which could be done with phase-modulated light as in [197].

Through coarse tuning, we found the standard operating parameters to be a crystal temperature of about 33.3 °C and a pump wavelength of 404.2 nm, resulting in an idler at 822.5 nm.

Once triple resonance had been established, it was possible to reach the desired transition (^{87}Rb D_1 $F = 1 \rightarrow F' = 1$ for the hyperfine and $F = 2 \rightarrow F' = 1$ for the Zeeman memory) over several weeks with only minor adjustments of the temperature and strain. On longer timescales, the desired signal frequency can not be reached anymore through the standard fine-tuning procedure. Triple-resonance conditions can then not be met anymore at the same pump frequency. When this happens, visible deterioration of the pump cavity modes can be recognized, hinting at pump-induced damage. There are different possible damage processes that can occur in KTP [198, 199]. Some of these, as the notorious gray tracing, are easily recognizable, while others, as photorefractive damage, stay invisible to the eye. Our damaged crystal had no visible defects. Heating up the crystal (also referred to as thermal annealing) has been shown [200, 201, 202] to reverse and even prevent these types of damage in crystals. Our few attempts to recover our crystals, however, remained unsuccessful.

A more in-depth investigation of this effect was beyond the scope of this work, especially considering the fact that this deterioration was not observed in the later version of the single-photon source. These kinds of damage mechanisms are localized to a specific pump mode. When a particular pump mode was deteriorated, we could resume operation by tuning the pump frequency by a few hundreds of GHz and adjusting the crystal temperature accordingly in order to find another set of parameters for which the phase-matching and triple-resonance conditions for the given signal target frequency were met. After readjustment, a similar performance could be established. Towards the end of the crystal's lifespan we also pumped off-resonantly. Even though the achieved emission rates were lower, this way of operation saved us much effort of trying in vain to lock on the deteriorated pump modes.

3.1.3 Characterization of the Optical Output

We start by characterizing the OPO cavity properties. Figure 3.4 shows the DFG peak (blue) of a 7 mm crystal recorded after an 8 mm etalon and the corresponding fit (red) yielding a $\delta\nu_{\mathrm{DRO}} = 195(1)$ MHz. Since we usually record our DFG signals

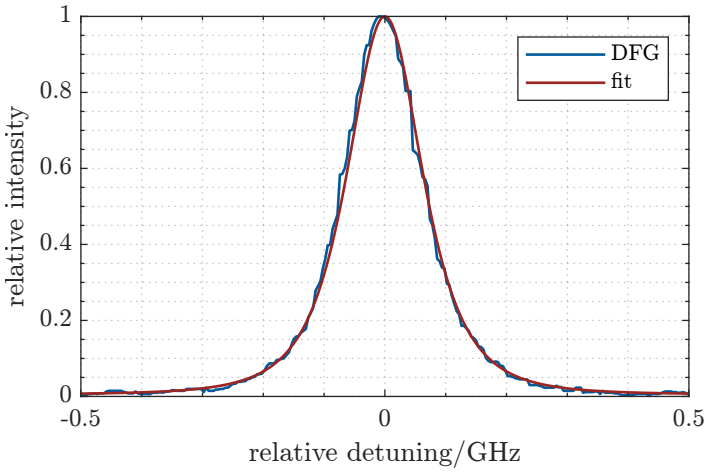

Figure 3.4: DFG signal (blue) induced in a 7 mm crystal measured after an 8 mm etalon. From a fit (red) the bandwidth $\delta\nu_{\mathrm{DRO}} = 195(1)$ MHz before filtration is extracted. The fit model is described in the main text. The modulation that can be seen on the blue trace originates from an external cavity that can form between the flat surface of the ppKTP crystal and the etalon.

L (mm)	FSR_s (GHz)	$\delta\nu_s$ (MHz)	$\delta\nu_{\mathrm{DRO}}$ (MHz)	$\mathcal{F}_s$
7.0(1)	11.22(14)	260.7(13)	195(1)	43(1)
5.0(1)	15.5(2)	424(2)	227.0(14)	37(1)
3.0(1)	27.4(4)	937(5)		30(1)

Table 3.3: Experimental data. The linewidth of the signal cavity $\delta\nu_s$ is measured with seed laser. $\delta\nu_{\mathrm{DRO}}$ obtained by measuring the DFG peak width (and correcting for etalon narrowing where necessary).

after the filtering stage, I fit the curve with the product of two Lorentzians in order to extract the $\delta\nu_{\mathrm{DRO}}$ linewidth of the idler photons. While the FWHM of one Lorentzian is a fit parameter, the second width is fixed to the FWHM bandwidth of the used etalon. The same x_0 is assumed for both Lorentzians in the fit. The bandwidth of the etalon is set to 274.4 MHz. Table 3.3 summarizes the OPO properties for the different crystals we operated the setup with. The FSR values are in good agreement with the expected ones. The cavity damping factor for the signal photons, however, deviates from the theoretical value. The measured value depends on the one hand on how well the dielectric coatings match their specifications and on the other hand on how well the seed laser is coupled into the OPO. Furthermore, losses within the medium can broaden the linewidth. The tabulated finesse is calculated from the FSR and cavity decay rate found in the table as well.

In order to characterize the quality of the optical output emitted by the source, we measure its second-order correlations. We perform time-correlated single-photon counting (TCSPC, or TTTR – time-tagged time-resolved) measurements. The single-photon detections are recorded as timestamps by a time-to-digital converter,

often referred to as 'time-tagger'. From these timestamps we compute arrival-time differences between the detectors. In order to obtain the second-order correlations, we perform start-multistop measurements. The idler photons are detected by a single SPAD while the signal photons are counted by two SPADs set up in HBT configuration. This allows us to compute the cross-correlation and the signal-signal autocorrelation, both conditioned on the idler and not, with a single dataset. To measure the autocorrelations for the idler we switch from the single SPAD to the HBT stage usually used for the signal.

To give an impression about the performance of the source, I give a specific example of generation rates measured with the 7 mm crystal. For a typical pump power of $P_{\text{pump}} = 1.2(1)\,\text{mW}$ we usually detect single photons in the signal and idler arms with rates of $c_s = 6.8 \times 10^4\,\text{cts/s}$ and $c_i = 1.7 \times 10^4\,\text{cts/s}$ respectively. The idler rate is significantly lower due to the stringent spectral filtering necessary to obtain only one remaining spectral mode. Integrating the total number of coincidences over an 8 ns window around the detection of the idler photon, we get a detected signal-idler rate of $r = 4.6 \times 10^3\,\text{pairs/s}$, or normalized to the pump power $r' = 3.8 \times 10^3\,\text{pairs/(s mW)}$. From these values we find the efficiencies $\eta_s = \frac{r}{c_i} = 27\,\%$ and $\eta_i = \frac{r}{c_s} = 6.7\,\%$. By correcting for the detector efficiency of $\eta_{\text{det}} = 60(6)\,\%$ we obtain a heralding efficiency of $\eta_h = \frac{r}{\eta_{\text{det}} c_i} = 45(5)\,\%$, expressing the likelihood of having a signal photon coupled into the signal fiber after having detected an idler. The heralding efficiency seems to be limited by the mode-matching to the optical fiber and by the quality of the reflective coatings. As we will see later in the double-resonant scenario, more photons than one would expect from the design specifications exit the cavity on the wrong side. The photon pair generation rate inside the OPO can be expressed as the measured pair detection rate r corrected for the losses and can conveniently be rewritten in terms of the measured counts $R = r/(\eta_s \eta_i) = c_i c_s / r$. Normalizing the internal generation rate to the pump power P_{pump} yields $R' = c_i c_s/(r P_{\text{pump}}) = 2.1(2) \times 10^5\,\text{pairs/(s mW)}$. The uncertainty follows from the uncertainty in the determination of the pump power.

For the following measurements the source was operated off-resonantly. Figure 3.5 shows the second-order correlations used to characterize the quality of the emitted photon pairs. Panel (a) shows the time-resolved cross-correlation in 162 ns bins. The plotted data set was recorded at a measured herald rate of $1.44 \times 10^4\,\text{s}^{-1}$. The signal-idler cross-correlation shows super-thermal bunching at zero time delay, i.e. $g^{(2)}_{s,i} \gg 2$. This is a first indicator for non-classical light. The cross-correlation is modeled as an asymmetric bilateral exponential function [109]. Due to the different filtration and slightly different bandwidths of the emitted photons, the coherence time corresponds to the cavity damping rate and thus differs for the signal and idler photons. I fit the data in panel (a) with a convolution of the model and the instrument response function (IRF) of our detector (detailed description follows)

$$g^{(2)}_{s,i\,\text{fit}}(\tau) \propto \left(\text{IRF} * \begin{cases} \exp\left(\tau/\tau_i\right) & \text{if } \tau < 0 \\ \exp\left(-\tau/\tau_s\right) & \text{if } \tau \geq 0. \end{cases} \right)(\tau) \tag{3.4}$$

The amplitude for the fit is not modeled but fixed to the maximum value of the

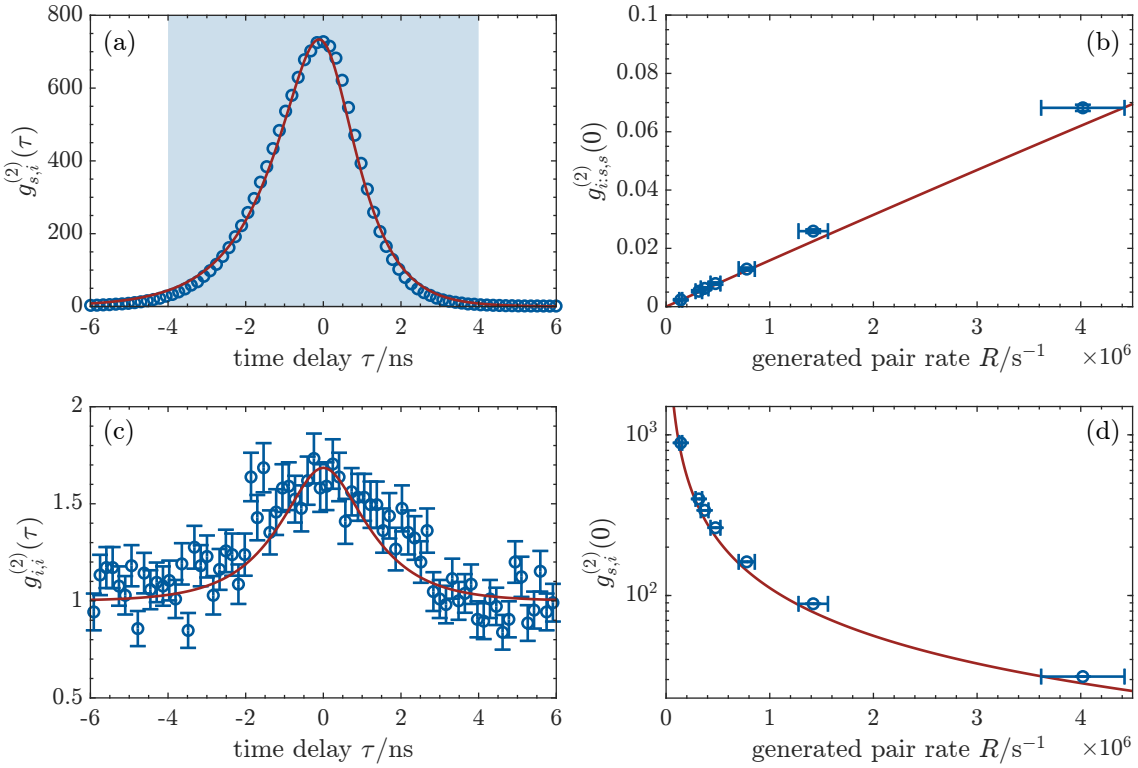

Figure 3.5: Second-order correlations of the 7 mm long OPO. All the plotted models (red) are explained in the main text. **(a)** Time resolved signal-idler cross-correlation in 162 ps bins measured at a herald rate of $1.44 \times 10^4\,\mathrm{s}^{-1}$. The shaded region corresponds to the 98×81 ps ≈ 8 ns coincidence window used to evaluate the points in (d). The shown data set corresponds to the fourth point in the bottom right plot. The asymmetry in the cross-correlation peak is expected, since signal and idler photons undergo differently stringent spectral filtration. The solid red line represents a fit with Eq. (3.4), used to extract the cavity damping rate of signal and idler. **(b)** Conditioned signal autocorrelation evaluated for one single bin of 8 ns as a function of the pair generation rate. In red, the theoretical model with no free parameters is shown. We measure conditioned signal autocorrelations $g^{(2)}_{i:s,s}(0) < 0.01$ for generation rates up to 5×10^5 pairs/s, expressing the single-photon nature of the heralded emission. For higher generation rates, corresponding to stronger pumping, the multi-pair generation probability increases. **(c)** Time resolved (unconditioned) idler-idler autocorrelation with 162 ps bins resolution. The model (red) takes account of the detector response function, and is in good accordance with a 'true' peak value $g^{(2)}_{i,i}(0) \approx 2$, corresponding to a single mode. **(d)** Signal-idler cross correlation as a function of the pair generation rate for an 8 ns coincidence window, integrating over the whole peak. This corresponds effectively to a time average. The theoretic model and its relation to (b) is explained in the main text.

data. The cavity damping rates obtained from the fit are $\tau_s = 212(4)\,\mathrm{MHz}$ and $\tau_i = 135(2)\,\mathrm{MHz}$. These rates are slightly narrower but still compatible with the measured FWHM of the seed $\delta\nu_s$ and the width we would expect for the idler after the etalon. The latter is approximated by calculating the width of the product of two Lorentzians, one with $\delta\nu_s$ and the second with the width of the etalon, which results in a bandwidth of 153 MHz. The discrepancy can be explained with an imperfect incoupling of the seed in the measurement, which leads to broader cavity resonances. The shaded area in the figure represents the 8 ns coincidence window used for the two panels on the right-hand side.

Through the finite time resolution of detector and the minimum bin width of the time-tagger, the conditioned autocorrelation at zero time delay $g^{(2)}_{i:s,s}(0)$ is technically inaccessible. The conditioned autocorrelation necessarily corresponds to the time-average over a specific coincidence window of the normalized Glauber autocorrelation [89]. The $g^{(2)}_{i:s,s}(0)$ gives the average single photon purity of the heralded single photon within the chosen window. The reported results will hence depend on the choice of this time window. This consideration is true for other coherence functions as well. The impact of this 'freedom of choice' is particularly great for the conditioned autocorrelation, since being a measure for the "state accuracy" it is a metric often used for comparing different work. To allow for fair comparisons it is important to know the considered time window. The choice of the duration of the ROI is a trade-off between high heralding efficiencies and low values of $g^{(2)}_{i:s,s}(0)$. We chose to consider a coincidence window of 8 ns having the storage and retrieval experiments in mind. The heralding efficiency is too important for the interfacing of the source with a memory – we need to account for all the photons in the source output, even if this lowers the $g^{(2)}_{i:s,s}(0)$ value. Using such a long coincidence window makes the correction for jitter unnecessary, since its effect is captured in the time-average.

After some fruitful discussions with Pavel Sekatski, we model the data with equation (24) he derives in [203] for non-number-resolving detectors. This theory is expressed in terms of the photon-pair generation probability p, which for continuous wave (CW) pumping and in the regime far below threshold can be approximated to be $p \approx R\Delta t$ [204, 205]. This approximation is valid for the regime far below threshold, i.e. for $p \ll 1$, which makes it less troublesome dealing with an unbound probability. The model from [203] also allows us to account for dark counts and detector inefficiencies. The dark counts of the SPADs used for this experiment are in the order of $10\,\mathrm{s}^{-1}$, which, being a factor 100 smaller than the lowest emission rate characterized, we deem as negligible. In Ref. [206] a visualization of the effect of high dark count rates on the conditioned autocorrelation can be found. Furthermore, we already corrected for the inefficiencies when deriving the generation rate R up for deriving p. The expression consequently simplifies to

$$g^{(2)}_{i:s,s} = 2p - p^2 \ . \tag{3.5}$$

Figure 3.5(b) shows the good agreement of the measured conditioned signal autocorrelation for increasing generation rates with the model, which has no free parameters. For pair generation rates of up to 5×10^5 pairs/s we measure conditioned

signal autocorrelations $g^{(2)}_{i:s,s}(0) < 0.01$. This indicates that multi-photon generation is strongly suppressed. The cut-off value is somewhat arbitrary, but allows for comparison between systems.

In the limit of operation far below threshold, the unconditioned autocorrelation is a good estimator for the number of modes N, with $g^{(2)}_{x,x}(0) = 1 + 1/N$ [207]. For measuring this quantity, a herald rate of $1.7 \times 10^5\,\mathrm{s}^{-1}$ is chosen as a compromise between a low multi-photon generation probability and a large number of coincidences, which directly influences the measurement time. For 142 min of integration, still, only about 7000 coincidences are measured within a coincidence window of 8 ns. This causes fairly large error bars for the autocorrelation time-resolved in 162 ps bins. Since the idler photons are spectrally filtered, their linewidth is determined by the product of the 195(1) MHz linewidth of the source cavity and the 274(4) MHz line of the 8 mm idler filter, both being Lorentzians. The shape of the ideal autocorrelation function is given by [89]

$$g^{(2)}_{x,x}(\tau) = 1 + \exp\left(-2\frac{|\tau|}{\tau_0}\right). \tag{3.6}$$

For Lorentzian light, linewidth and coherence time relate through the expression $\tau_0 = (\pi\Gamma)^{-1}$ [208]. In our scenario $\tau_0 \approx 2.2$ ns. Since this value is only reached at zero delay, the measured value will be lower, even for a single mode, in case the detector's timing jitter is comparable to τ_0. The model plotted in Fig. 3.5(c) is the convolution of the instrument response of our HBT setup and the ideal autocorrelation. The timing jitter of our HBT setup is best modeled as a hyperbolic secant with a FWHM= 1.09(5) ns, as is shown later in Section 3.1.6. The good agreement of the model with the data is an indicator for the contamination of unwanted modes being negligible after filtration. For a single 8 ns bin we measure $g^{(2)}_{i,i} = 1.338(16)$. However, for the signal photons, which are not significantly filtered, other modes are present and no features can be seen in the autocorrelation. We measure $g^{(2)}_{s,s}(\tau) \approx 1$ everywhere, within the measurement uncertainty. This confirms that filtering the herald is indeed necessary for single-mode emission of the heralded photon source.

Figure 3.5(d) shows the trend of the cross-correlation with the pair generation rate. From the model in Eq. (3.5) an expression for the cross-correlation can be derived by using the following relation between second-order coherence functions derived from Bayes' theorem [209]

$$g^{(2)}_{i:s,s} = \frac{g^{(2)}_{s,s}\, g^{(2)}_{i,i}}{g^{(2)}_{s,i}}. \tag{3.7}$$

We account for the contamination of unfiltered modes by using the measured value of the idler autocorrelation for the signal $g^{(2)}_{s,s}$ as well. The assumption made here is reasonable, since the heralded signal, as it appears in the cross-correlation and the heralded signal autocorrelation, inherits the mode purity of the idler, since it is conditioned on the latter.

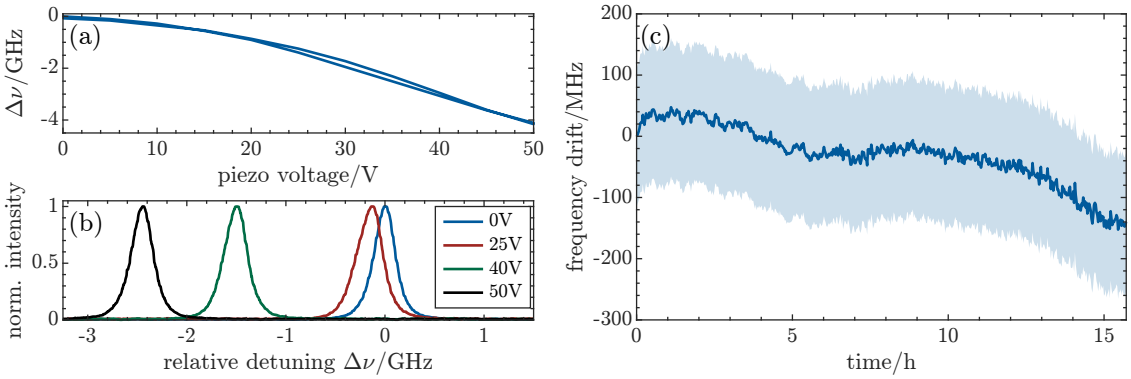

Figure 3.6: Strain tuning behavior and stability of the source. The shown data were acquired in Berlin with the 5 mm crystal. **(a)** Strain tuning of the seed cavity resonance. At about 20 V the piezo makes full contact and the tuning follows a linear behavior. **(b)** Tuning of the DFG signal. Since the applied strain changes the phase-matching conditions the crystal temperature needs to be slightly adjusted. This change, however, counteracts the strain tuning, making the achievable range smaller than what we see in (a). The DFG peak amplitudes are normalized. **(c)** Long term stability of the DFG frequency. For about 16 h the seed laser is scanned around the signal frequency while the crystal is pumped. The seed frequency corresponding to the maximum of the DFG peak is recorded. The shaded area represents the FWHM of the DFG peak. An average drift of about 10 MHz h^{-1} is measured. The system is free-running: there is no active feedback on the signal nor idler frequency, merely the pump laser is locked to the OPO.

3.1.4 Strain Tuning

As mentioned earlier, the monolithic cavity design comes with advantages in heralding efficiency and stability. However, it also comes with the loss of available degrees of freedom with respect to an external cavity, composed of independent optical elements. In order to still be able to meet the triple-resonance conditions, we apply mechanical strain with a piezoelectric actuator to the non-linear crystal. The strain deforms the refractive index ellipsoid through the elasto-optic effect, allowing for fine-tuning of the triple-resonance conditions [210, 211].

The strain is applied orthogonally to the direction of propagation of the light. A fine-threaded screw is used to pre-strain the piezoelectric actuator. For this measurement the piezo was pre-strained to make full contact with the copper plate pressing on the KTP crystal at approximately 20 V. From there on we observe a nearly linear tuning behavior (see Fig. 3.6(a)). Some hysteresis is observed with this method, as can generally be expected from piezos. By frequency-tuning the seed laser, by applying a voltage to the piezo, and subsequently readjusting the crystal temperature by a few tens of mK, it is possible to establish triple resonance. The temperature adjustment also affects the phase-matching conditions. Since the phase-matching bandwidth is several hundreds of GHz broad, this only has a minor effect. The amplitude of the DFG peak can, however, decrease by up to 50 %. The necessary temperature adjustment, however, acts in the opposite direction of the strain-tuning, effectively

reducing the maximal tuning range of the DFG peak, compared to the seed cavity resonance. The tuning range of the seed frequency for which we observe DFG with this technique is $> 2\,\mathrm{GHz}$, as can be seen in Fig. 3.6(b). This tuning range was not always sufficient to tune the resonance to the desired signal frequency. In these cases the frequency tuning process needed to be started again with a different set of parameters.

3.1.5 Long-Term Stability

To assess the long-term frequency stability of the OPO the source was operated continuously for 16 h while being seeded by the seed laser. The laser frequency was repeatedly scanned around the $^{87}\mathrm{Rb}$ D_1 $F = 1 \rightarrow F' = 1$ transition and the seed frequency corresponding to the peak of the DFG signal was recorded. The measured frequencies are plotted in Fig. 3.6. The frequency of the seed laser was initially calibrated through an atomic spectroscopy in order to get an absolute frequency reference for the scan range. While the pump laser was locked to the OPO, no active feedback to stabilize the output frequencies was used. Over the course of the measurement an average frequency drift, defined as the standard deviation σ divided by the measurement time Δt, of about $\delta\nu = 3\,\mathrm{MHz\,h^{-1}}$ is reached.

3.1.6 Further Characterizations

As a conclusion to the description of the triple-resonant OPO I briefly describe some additional characterizations of the system that we performed.

Herald Filter Stage In order to assess the suppression of unwanted modes we characterized the herald filter stage. For this purpose we used a widely tunable CW titanium sapphire (Ti:Sapph, M Squared SolsTiS) laser that we tuned to the idler wavelength. The filter stage was characterized with both filters in place, so that the setup did not need to be altered.

The Ti:Sapph is scanned over a $\sim 6\,\mathrm{GHz}$ range in order to measure one transmission resonance of the etalon with the amplified photodiode $\mathrm{PD}_{\mathrm{DFG}}$. Different scans are recorded for wavelengths in a range of $\pm 2\,\mathrm{nm}$ around the peak transmission wavelength of the filter stage. The power of the laser is monitored before the filter system, allowing for later correction of wavelength dependent power variation. In Fig. 3.7(a) the maximum transmission of each single measurement scan, normalized to the input power, is shown as a function of the central wavelength of the Ti:Sapph scans. A Gaussian fit yields a FWHM $= 0.57(5)\,\mathrm{nm}$. The dynamic range of a single gain setting of the amplified photodiode is not large enough to capture the whole extinction ratio. We therefore need to increase the gain the further we move away from the idler frequency. From Fig. 3.7(a) we suspect that the measured extinction is strongly dependent on the gain, manifesting as plateaus in the logarithmic

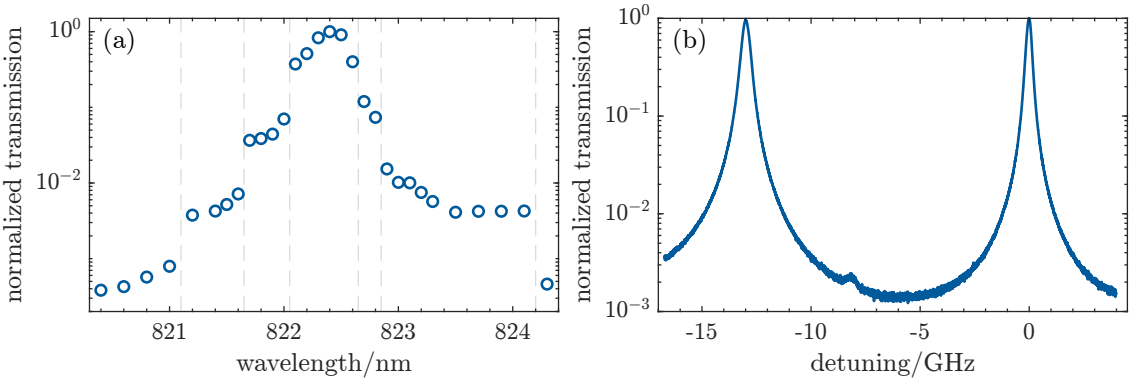

Figure 3.7: Characterization of the herald filter stage. **(a)** Transmission of the IF. The Ti:Sapph is scanned to measure the peak transmission through the etalon at different wavelengths over a range of ∼4 nm. By plotting the maximum transmission of each scan the IF transmission is reconstructed. The artifacts/steps in the data originate from different gain settings of the photodiode used for this measurement. The dashed vertical lines represent when the gain was switched. **(b)** Transmission through the 8 mm etalon close to the resonance frequency of the IF. The Ti:Sapph was attenuated to few photon level in order to perform the measurement with a SPAD for increasing the dynamic range with respect to the amplified photodiode. The slight asymmetry in the linewidths could be caused by a non perfectly linear frequency scan of the laser. The small feature visible at −8 GHz is most likely a higher order mode of the etalon. The plot is cropped to one of the scan flanks.

plot. The extinction ratio at the tails (>1 FWHM away from the maximum) can be conservatively estimated as $<10^{-2}$.

For the etalon characterization we again used the Ti:Sapph, this time attenuated to the few photon level (nominally 14 orders of magnitude) before the etalon, so that we can use a SPAD as detector and achieve the necessary dynamic range. The wavelength was tuned to 822.4 nm to match the maximal transmission through the IF. We now scan the Ti:Sapph over a range of 24 GHz allowing us to measure two etalon resonances around the IF's maximum transmission. Figure 3.7(b) shows the recorded resonances. The FSR, calculated form the etalon thickness, defines the frequency scale, which is in good agreement with the laser scan settings. A FWHM = 274(4) MHz is extracted from the data and an extinction ratio of $\sim 1.2 \times 10^{-3}$ at FSR/2 is measured.

It is crucial for this characterization that the etalon is properly aligned, else the measured linewidth becomes broader. We optimized the filter stage right before this measurement and recorded an overall peak transmission through the system of about 86 %.

Measuring the HBT Detection Jitter The SPADs (Excelitas SPCM-AQRH-16) we used for the previously reported measurements are specified to have a minimum response time of 350 ps. Using two such detectors in HBT configuration results in a timing resolution of about 500 ps, since jitter adds quadratically. Using

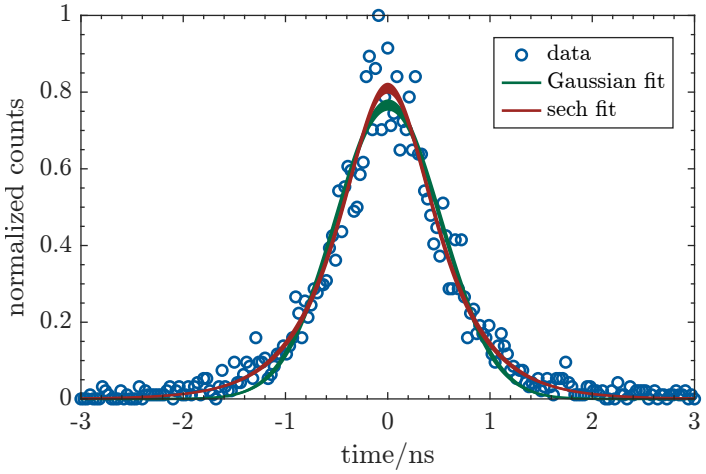

Figure 3.8: Determining the timing jitter of the HBT setup as it was used in Section 3.1. The measured response is obtained by sending ps laser pulses to the HBT setup. The FWHM of the Gaussian and the *sech* fit are 1.21(5) ns and 1.09(5) ns respectively. The uncertainties include a bin width of 30 ps as an estimate for systematic errors. The *sech* fit models the tails of the response better. The thickness of the lines of the fit models represents their 95 % confidence interval.

this value to model the data in Fig. 3.5 results in a poor agreement. Non-optimal alignment of the detection stage increases the jitter. Unfortunately, for the measurements presented in Section 3.1, this was the case. For instance, only one lens was used to focus the beams onto the two SPADs. In order to be able to satisfyingly model the measured data it was necessary to empirically determine the timing jitter of our detection stage.

When we perform a measurement with the HBT setup, the result will be the convolution of our input pulse with the instrument response function. If the pulse is short enough compared to the time scale of the timing jitter, it can be approximated as a Dirac delta distribution. Since the delta distribution is the identity element of the convolution, through this measurement we gain access to the detector response. For this measurement Alisa Javadi and Natasha Tomm kindly shared some light from their pulsed Ti:Sapph (Coherent Mira 900) with us. The laser periodically emits 12 ps pulses at 820 nm which are attenuated and subsequently sent to the HBT setup. A fast time-tagger (Swabian Instruments) records the time difference between the two detectors. The resulting instrument response is shown in Fig. 3.8.

The temporal response of SPADs exhibits fast transients in the peak followed by exponential behavior [212]. While the FWHM of the measured response changes only by little with the choice of the model function, a hyperbolic secant matches the tails better than a Gaussian. A *sech* fit results in a FWHM of 1.09(5) ns.

After this characterization was performed, we realigned the HBT stage adding more degrees of freedom. The timing jitter should now be closer to the specified values.

3.2 Double-Resonant OPO

Even though the performance reached by the triple-resonant single-photon source design met the set requirements for an interfacing attempt with the quantum memory, it showed some major flaws in its operability. Over time it became more and more difficult to satisfy the triple-resonance conditions. Deterioration of the pump cavity modes required regular realignment with a different set of parameters (pump frequency, crystal temperature, strain), up to the point where the crystals were so damaged that it would become more convenient to operate the OPO off-resonantly. In the one and a half years we worked with the triple-resonant design two crystals broke irrevocably, one of which in a pretty dramatic way: one of the end-facets chipped off.

We revised the design of the source with a focus on the ease of operation and longevity of the crystals . By changing to a double-resonant mode of operation with a double-pass pump configuration we can freely tune the pump frequency, making it a truly independent degree of freedom. The low finesse cavity for the pump we had in the previous design led to a ratio of circulating to input power of less than 7. Using a diode laser, which in our case outputs up to about 100 mW, and the double-pass scheme, we have enough power at our disposal even without the cavity. The limiting factors for the achievable generation rates will be the single-photon quality and the thermal stability. With the pump frequency as new freely adjustable parameter, we can abandon the strain tuning. This relaxes the requirements to the crystal's and its oven's surfaces, on which we previously applied strain, since they don't need to be highly polished anymore. The new crystal copper mount features a notch that allows parallel alignment of the crystal to the oven wall. This feature allows us in principle to swap the crystal without having to restart the optical alignment procedure from scratch. In the previous design, the crystal orientation did not necessarily correspond to the one of the oven, making the alignment procedure counterintuitive. For the next design iteration it would be even more favorable to rotate the orientation of the crystal within the mount by 90°. Signal and idler polarization would then directly coincide with the desired ones making the half wave-plate after the source superfluous.

Since with the previous implementation we cycled through several crystals, we decided to use a double-stage TEC (Thorlabs TECD2, $T_{max}^{\text{hot}} = 200\,^{\circ}\text{C}$, $\Delta T_{max} = 108\,^{\circ}\text{C}$) and integrate two resisitive cartridge heaters (Thorlabs HT15W) for the possibility of *in situ* thermal annealing, which allegedly helps reversing and even preventing damage to the crystal. Furthermore, the oven was designed to offer the possibility to connect a hose for flooding it with gas, since particular atmospheres (e.g. N_2) allegedly can help in the recovery process of the crystal[1]. These new

[1] Apparently 'dry' atmospheres can help with the recovery of hygroscopic crystals. The reaction of these type of crystals with the humidity in air can also lead to damages to optical coatings (one could say: they are still 'alive' compared to standard optical substrates, according to the coating company). In hindsight, however, I'm not sure how helpful this would be in our case since KTP is not hygroscopic.

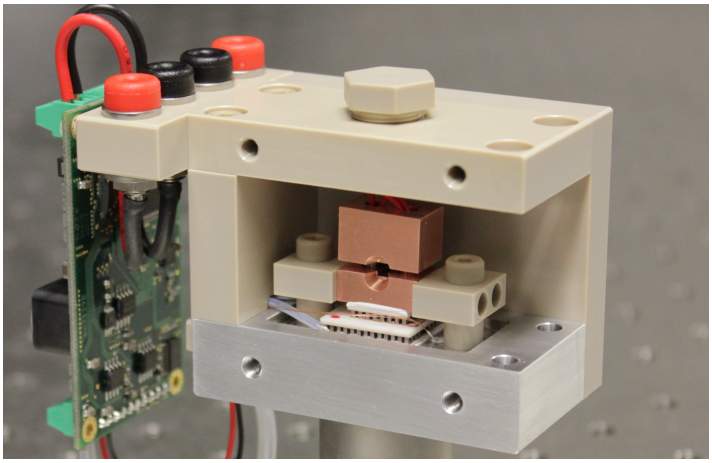

Figure 3.9: Photograph of the oven design for the double-resonant OPO with the housing partially removed. The crystal is enclosed in a copper holder with a notch for coarse alignment of the crystal's orientation. The ppKTP temperature is set by a double-stage Peltier and stabilized with a TEC-controller. The bottom of the enclosure made out of aluminum acts as a heat sink, while the walls and the top are made out of PEEK for better thermal insulation.

features, however, remain untested since the double-resonant OPO has now been operated for over two and a half years with the same crystal without any noticeable signs of deterioration. A further small modification of the design was to change the wall material of the oven to polyether ether ketone (PEEK) for better thermal insulation, while the base plate of the oven was kept out of aluminum to serve as a heat sink to the TEC. In the double-resonant implementation of the source, the pump laser is also first fiber coupled into an end-capped PM fiber for spatial mode cleaning before it reaches the crystal.

By operating the OPO in a double-pass pump configuration we lose the ability to lock the frequency of the laser to the monolithic cavity. To make sure that the pumping frequency stays stable over time, we set up an external reference cavity. A detailed description of the frequency lock follows in Section 3.2.2.

For the double-resonant OPO we had to commission new monolithic cavities. Two 5 mm long pieces of ppKTP were left from the original growth process. The plane facets of the crystals were coated in the same run as the crystals used for the triple-resonant OPO. The curved surface was polished in a new run to the same specs by the same manufacturer, while the coatings were applied by *Laseroptik* with the nominal reflectivities of $R^c_{s/i} > 99.9\,\%$ at $810\,\text{nm} \pm 15\,\text{nm}$ and $R^c_p < 1\,\%$ at 405 nm. We used this opportunity to send in the 5 mm crystal with the shattered coating on the curved surface to be repolished and recoated, now for double-resonant operation, as well.

Tuning the Output Frequency With the absence of pump resonances, the procedure used to tune the source to the desired frequency had to be adapted. As

before, we start by coupling the seed laser to the OPO and fine-tune the crystal temperature so that a seed cavity resonance coincides with the desired output frequency. Instead of keeping the pump frequency fixed at one of its resonances while scanning the seed laser, as before, the new tuning routine for finding the DFG signal consists of scanning the pump laser while the seed is free-running at the desired output frequency. Once a good DFG signal is generated, the pump laser can be locked and the seed laser disconnected from the setup. The source reliably generates photon pairs for up to a day. Once the source has been tuned to a specific frequency, the daily routine when turning it on consists merely of getting the pump laser to emit at the right frequency, checking that the cavity resonance is still at the target frequency, and adjusting the lock parameters. The double-resonant source can usually be operated for months before parameters like the crystal temperature or the pump wavelength need to be adjusted. The beam pointing of the pump laser needs more regular adjustments to maximize the DFG amplitude, though.

3.2.1 Characterization of the Optical Output

The assembly and characterization of the double-resonant OPO is thoroughly described in BC's master's thesis [182]. I will therefore limit myself here to describe a few selected results. The performance of this source was overall equal or better than the triple-resonant implementation. A summarized version of the used parameters and achieved results can be found in Tab. 3.4. The characterization measurements were performed analogously to those previously presented for the triple-resonant case.

We noticed that the FSR of the double-resonant cavity is larger than what we measured for a crystal of the same nominal length in the triple-resonant case. It turns out that one of the crystals we sent for polishing and coating on the second run was actually only 4.7(1) mm long before the procedure, being the 'leftover' piece from the original cutting. From the measured signal FSR we derive that the monolithic cavity that we characterized has a length of 4.5(2) mm, giving us a good reason to believe that we used this shorter crystal for the double-resonant OPO, which shrunk even further due to material ablation during the polishing procedure.

A measurement with a single-photon resolving spectrometer shows that the emission of the double-resonant OPO is cleaner overall than what we saw in Fig. 3.3(a). This purer spectrum with fewer modes leads to a favorable comparison of rates even with less pump power, since a higher fraction of photons are emitted in the desired mode. For low pumping powers even a bunching feature in the unconditioned autocorrelation of the unfiltered signal can be detected. For 1 mW of pump power a $g^{(2)}_{s,s}(0) > 1.3$ was measured for a time resolution of 162 ps. This was not possible with the old design, where no such feature could be measured.

Inspired by one of the referees of [64], we decided, given that we designed the new setup to have enough space to perform this measurement, that it would be interesting to measure the correlation of the backwards emitted photons. A distinct correlation

Symbol	Property	Value
L	Length, after polishing	4.5(2) mm
	Cross-section	$1 \times 2\,\mathrm{mm}^2$
$\Delta\nu_{\mathrm{FSR}}$	FSR	17.5(2) GHz
$\mathcal{F}$	Measured finesse	41(1)
$\delta\nu_s$	Cavity damping factor (signal)	426(2) MHz
$\delta\nu_{\mathrm{DRO}}$	Bandwidth	373(1) MHz
η_h	Heralding efficiency	53(5) %
Δt	Standard coincidence window	6.48 ns (80×81 ps)
	Maximum pair rate for $g^{(2)}_{i:s,s} < 0.01$	$7.5(8) \times 10^5$ pairs/s
$g^{(2)}_{i,i}$	Idler autocorrelation	1.28(5)
$\mathrm{P}^{typ}_{\mathrm{pump}}$	Typical pumping power	4.5 mW
r^{typ}	Typical herald rate	$1.5 \times 10^5\,\mathrm{s}^{-1}$
$g^{(2),typ}_{i:s,s}$	Typical conditioned signal autocorrelation	0.04
$\eta_{h,\mathrm{DL}}$	Heralding efficiency after fiber and with OI	40(4) %
$\tau_{s,\mathrm{min}}$	Minimum storage time when switching the source	130 ns

Table 3.4: Summary of the properties of the monolithic double-resonant OPO as specified in [182]. In the lower half of the table the parameters used for the interfaced measurements with the Zeeman memory are listed.

peak between the backwards emitted signal photons and the forward emitted heralds can be measured. BC estimates the fraction of signal photons leaving the cavity on the wrong side to be approximately $\eta_{\mathrm{back}} = 13\,\%$, which is considerably higher than expected. A possible cause could be the quality of the coating on the curved surface achieved during manufacturing, since crystals apparently are harder to handle than standard substrates for optics. JW and his group are currently trying to reduce the backward losses in a similar system by using the plane surface as rear mirror. Considering the measured heralding efficiency, a lower bound for the mode-matching into single-mode fiber can be estimated to be around $\eta_{\mathrm{fiber}} > \eta_{\mathrm{herald}}/(1-\eta_{\mathrm{back}}) \approx 61.6\,\%$. Which, considering the transverse mode profile of the DFG, seems rather reasonable.

3.2.2 Locking the Pump Laser

A double-resonant OPO design with only a double-pass configuration for the pump beam does not result in a resonance. In order to still be able to stabilize the frequency of the pump laser we opt for a Pound-Drever-Hall (PDH) [213, 214] sideband offset lock on a passively stable external reference cavity. In brief, this technique relies, as

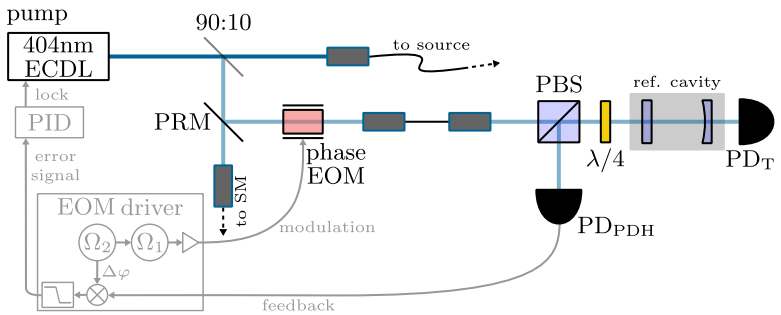

Figure 3.10: Setup of the 404 nm pump laser preparation stage with PDH sideband offset lock. While the spatial pump mode is cleaned with a PM fiber, a fraction of the light is picked off and sent to the PDH sideband offset lock. By locking onto the sidebands we are able to achieve a stabilized but tunable frequency. PD_T is used to facilitate the coupling of the cavity. A fiber-coupled spectrometer is used to monitor the pump frequency. Used abbreviations: 90:10 – beam sampler with the specified branching ratio; ECDL – external cavity diode laser; PRM – partially reflective mirror; EOM – electro-optic modulator; PBS – polarizing beam splitter; PD – photodiode; SM – spectrometer.

the PDH method, on the interference of the off-resonant sidebands which are fully reflected by the cavity and the phase-shifted 'carrier' component. This phase shift is strongly frequency dependent in proximity of the resonance and can be used to generate an antisymmetric error signal. In the sideband offset lock, instead of the carrier, a sideband is tuned into resonance with the cavity and is interfered with off-resonant sub-sidebands. The technique we implement is described in [215] as electronic sideband locking. The light at the carrier frequency ω_c is phase-modulated with an electro-optic modulator (EOM) driven with a signal which is itself phase-modulated. The drive signal has the carrier frequency Ω_1 and is modulated at Ω_2, with $\Omega_1 > \Omega_2$. The electric field of the light after the EOM can be expressed as

$$E = E_0 \exp\left[i(\omega_c t + \beta_1 \sin(\Omega_1 t + \beta_2 \sin \Omega_2 t))\right] \ , \tag{3.8}$$

where β_i is the corresponding modulation depth. By expanding this expression to the first order in β_i, it can be seen that the resulting spectrum contains sidebands at angular frequencies $\omega_c \pm \Omega_1$, as well as sub-sidebands at angular frequencies $\omega_c + \Omega_1 \pm \Omega_2$ and $\omega_c - \Omega_1 \pm \Omega_2$ [215]. The spectral structure around $\omega_c + \Omega_1$ with sidebands offset by $\pm\Omega_2$ is analogous to the standard PDH modulation spectrum (analogously for $\omega_c - \Omega_1$). If the modulation frequency Ω_1 is tunable within a range of FSR/2 it is always possible to tune one of the sidebands at $\omega_c \pm \Omega_1$ to the cavity resonance. Thus, with this method a stable but tunable frequency output can be achieved.

A sketch of the experimental setup is shown in Fig. 3.10. We use a temperature-stabilized phase-EOM (QUBIG PM9-VIS_1.5) with tunable resonance frequency to modulate sidebands onto the pump light. An amplified RF signal generator (QUBIG QDG10_A0.05-3.0W3) that can be tuned over a range of $\Omega_1 = 750 - 1500$ MHz, corresponding to half an FSR of our cavity, is used to generate the tunable sidebands

at an arbitrary position within the cavity's FSR. The driver system of the modulator furthermore includes a local oscillator (0.5 – 200 MHz) with precision phase control (< mrad) that is used to modulate the driving signal and later on for the internal heterodyne detection to generate the PDH error signal. The modulated light is then sent to the cavity. Our reference cavity has a hemispherical design where the curved mirror has a radius of curvature of −200 mm. For this reference cavity we use off-the-shelf mirrors with polished backside from *Thorlabs*. The mirrors are cemented with an optical adhesive to a 10 cm long spacer made out of Zerodur[2] (fabricated by *Aachner Quarzglas Technologie Heinrich*), a lithium-aluminosilicate glass-ceramic with an ultra low thermal expansion coefficient. The resonator lies on an aluminum platform with a notch, to avoid mounting stress, resting on a Sorbothane sheet for vibration isolation inside a small vacuum chamber. The pressure inside of the tube is reduced to ∼20 mbar for partial acoustic isolation. Figure 3.11(a) shows the transmission spectrum of the cavity with a modulation frequency of $\Omega_1 = 800$ MHz generating the first order side bands. A measured finesse of $\mathcal{F} \approx 340(30)$ is reached. A PBS and a quarter-wave plate, which is passed twice, are used to access the reflected beam, which is then detected by a fast photodiode (QUBIG PD-100-VIS). The recorded signal is sent back to the EOM driver in order to generate the PDH error signal by demodulating it with Ω_2. Finally, the error signal is sent to a PID regulator integrated in the laser controller, for locking.

Once a good DFG signal is established, one of the main sidebands is tuned to coincide with the pump frequency generating the maximal DFG emission. Changing the modulation frequency Ω_1 of the phase EOM requires adjusting the potentiometer determining the resonance of its RF circuit. The steep edge of the generated error signal at this point should coincide with the DFG peak. The phase, the frequency, and the power of the modulation at Ω_2 can be fine-tuned for generating a clean error signal. Once the pump laser is locked, the seed laser can be turned off, and the source produces single-photon pairs at the desired frequency.

In the current setup, the lock parameters need to be re-adjusted on a daily basis before the laser can be locked. A possible reason for this behavior is a small leak in the vacuum chamber, which leads the pressure to increase over time affecting the cavity spectrum. To assess the mid-term stability of the lock we performed the same measurement as described previously in Subsection 3.1.5. This time the seed laser was simultaneously referenced to an atomic spectroscopy so that any drifts of the scan range could be compensated later on. The data were recorded after we tuned the source to the D_2 line. The stability measurement was limited by the particular ECDL we used as a seed, which was prone to mode-hop near the frequency of interest. Still, the characterization measurements remains stable (at least) an order of magnitude longer than the usual integration times. If the frequency lock is set up well, the laser stays locked even over night. The recorded frequency over time is shown in Fig. 3.11(b). The resulting average frequency drift $\delta\nu = 5.9\,\mathrm{MHz\,h^{-1}}$ is in the same order of magnitude as in the triple-resonant OPO.

[2]Technical details for Schott Zerodur available at https://www.schott.com/en-us/products/zerodur-p1000269/technical-details.

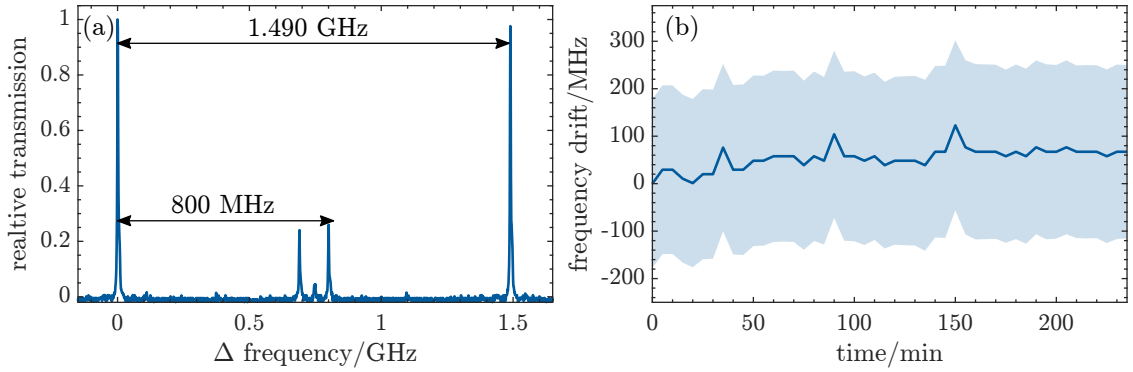

Figure 3.11: Characterization of the sideband offset lock. **(a)** Transmission spectrum of the reference cavity for a sideband modulation frequency of $\Omega_1 = 800\,\mathrm{MHz}$. Note that the sidebands lie crosswise with respect to their carrier. By using the modulation frequency as an absolute frequency reference we measure an $\mathrm{FSR} = 1.490(5)\,\mathrm{GHz}$. The sub-sidebands are too weak to be recognized. **(b)** Long-term stability of the double-resonant OPO while locked on the reference cavity. With the sideband lock we observe an average drift of $5.9\,\mathrm{MHz\,h^{-1}}$ of the seed frequency at which DFG is induced. The shaded area represents the FWHM of the DFG peak. This data measurement was performed with the source operating at the D_2 line.

3.3 Changes for Interfacing

During our first attempts to interface the photon-pair source as it was set up in [182] by BC with the Zeeman memory described in Ch. 6, we observed a constant noise floor of uncorrelated photons being emitted by the downconversion source as well as crosstalk between the two systems. Fortunately, we could suppress both effects with straightforward modifications. An experimental setup for the heralded photon source as it was used for the interfacing experiments is shown in Fig. 3.12. Details about the Zeeman memory will follow in Chapter 6.

Switching the Source In order to suppress the mentioned noise floor of resonant uncorrelated photons emitted by the source, we decided to switch off the source during the retrieval window of the memory. As an optical switch we use a fast Pockels cell (QUBIG PC2B-VIS) that responds to a follower TTL logic and rotates the linear polarization of the pump beam by 90°. Hereby, the type-II downconversion process is highly suppressed since the phase-matching conditions are no longer met. As the pump light remains incident on the OPO, no thermal drifts should be induced by the switching. The cell is driven by a high-voltage source (QUBIG HVS_1.0-1500) directly connected to it in order to minimize the capacitance ($< 10\,\mathrm{pF}$) and to allow for fast switching. For a periodic input signal the Pockels cell is supposed to reach a specified suppression ratio of 1:1000. Motivated by our low repetition rate we decided to mount the cell and driver assembly well thermally contacted onto an aluminum heatsink, although for MHz repetition rates it is recommended

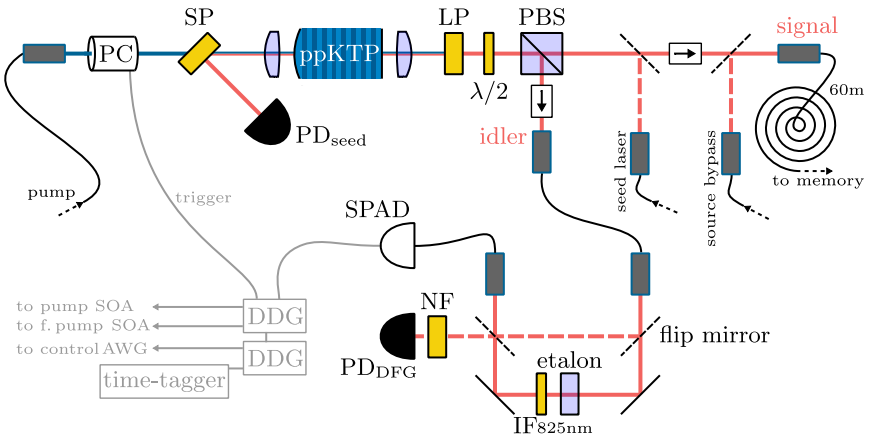

Figure 3.12: Experimental setup of the double-resonant source as used for interfacing with the ^{87}Rb D_1 line quantum memory. The pump laser is set up and locked according to Fig. 3.10. The crystal is pumped in a double-pass configuration. In order to suppress uncorrelated noise we switch off the source by rotating the pump polarization after the detection of an idler. Optical isolators reduce crosstalk between the systems. The detection of an idler triggers the electronic logic of the quantum memory. In the meantime the signal photon is sent through a delay line, consisting of a 60 m PM fiber. Flip mirrors allow for injecting the seed laser in counterpropagating direction and bypassing the source once it has been already interfaced with the memory. Used abbreviations: PC – Pockels cell; SP (LP) – optical short-(long-)pass filter; PBS – polarizing beam splitter; NF – 785 nm notch filter; IF – interference filter; PD – photodiode; SPAD – single-photon avalanche diode; DG – delay generator. For the operation at the D_2 line two minor adjustments need to be performed in the setup: A PBS needs to be added right after the PC in order to fully switch off the source and the IF in the idler filter stage needs to be removed.

to watercool it.

For a change, the integration of the Pockels cell into the setup proved to be a straightforward and uncomplicated task. The achievable suppression depends on the alignment of the beam passing the cell, the voltage, and the switching frequency. For the alignment, the optical technique known as conoscopy offers a systematic procedure that can be followed. By placing a 'bad' diffuser in front of the Pockels cell, part of the light is divergently scattered while the rest is transmitted. The polarization component orthogonal to the incident one is monitored on a screen. For a birefringent medium an interference pattern can be observed. For a uniaxial material, such as beta barium borate (β-BaB_2O_4 or BBO) which composes our Pockels cell, the interference pattern consists of a "Maltese"-cross (isogyre) and concentric rings (isochromes). Furthermore, the transmitted fraction of the incident beam can be recognized as a bright spot. Using the yaw, pitch, and Z axes of the stage the Pockels cell is mounted on, the pattern is centered onto the transmitted spot. The alignment procedure is voltage- and repetition frequency-dependent and might need realignment if these parameters are changed. For optimizing the extinction ratio at a certain frequency, the high-voltage control knob needs to be adjusted.

Figure 3.13(a) shows how the cell performs in a realistic scenario. The plot shows an arrival-time histogram of the transmission through the source-memory interface/assembly. The memory sequence is not started. The Pockels cell, however, is triggered by the detection of an idler photon. Zero time delay in this histogram corresponds to the time an idler is detected. The delayed signal is recorded at 260 ns. Upon detection of an idler photon, it takes about 130 ns for the source to switch off, which sets the minimal storage time of the memory if uncorrelated noise is to be avoided. The off-time of the cell is set to be 500 ns. The measured noise floor when the source is switched off is not dark-count limited.

Suppressing Crosstalk The uncorrelated noise floor emitted by the source is constant and switchable, which allows us to subtract it from the data. Turning on the quantum memory while the photon-pair source is being switched reveals some new features/structure in the background noise. We could eventually trace them back to be crosstalk between the lasers preparing the initial atomic state of the vapor in the memory and the non-linear crystal.

This new noise feature, see Fig. 3.13(b), shows a correlation with the atomic pump laser. There is a fixed delay between the switching time of the atomic pump laser and the noise feature, corresponding to an additional round-trip of the photons in the 60 m fiber used as delay line. The delay $\Delta t = 2 \times 294$ ns matches the measured fiber-induced delay of 285 ns plus an additional 9 ns. This small difference corresponds to roughly 3 m in free space and matches well the distances in the experimental setup: from the vapor cell to the signal collimator on the memory side and from the fiber coupler to the plane surface of the ppKTP crystal on the source side. The cell used in the memory experiment has wedged windows to suppress etalon effects and copropagating backreflections. However, the fluorescence induced by the

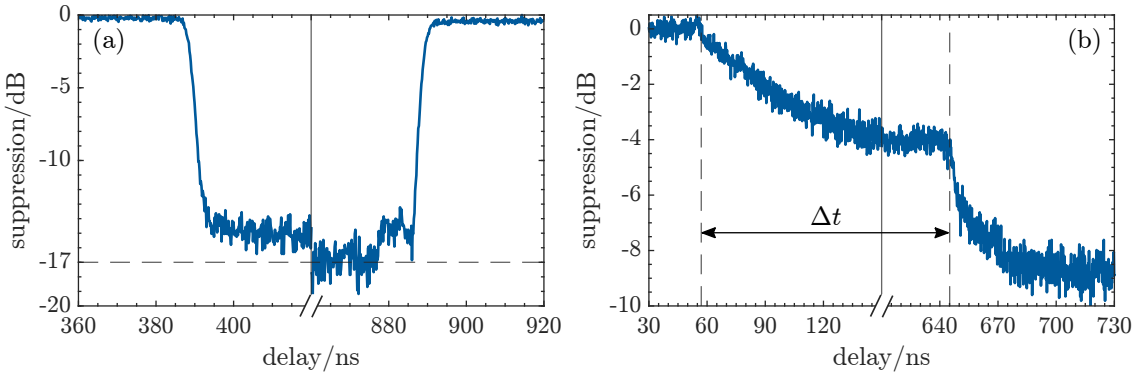

Figure 3.13: Arrival-time histograms characterizing the background noise when source and memory are interfaced. The zero time delay corresponds to the detection time of an idler. **(a)** Switching behavior of the Pockels cell. The uncorrelated noise floor can be suppressed by −17 dB. **(b)** Unless an optical isolator is placed in the signal arm we observe crosstalk between the two systems. Induced but switchable noise can make a round-trip from the memory to the source and back again delaying the complete switching by $\Delta t = 588\,\text{ns}$. Since the uncorrelated noise floor of the source is constant and switchable it was subtracted from the data outside of the switching region to exclusively show the effect of the crosstalk.

lasers pumping the atoms is emitted into a large solid angle. A fraction of these fluorescence photons is coupled backwards into the signal arm and reflected by the plane surface of the OPO, and thus detected a round-trip time later, after the atomic pump lasers have been switched off (for details about the switching see Chapter 6).

Placing an optical isolator (OI) in the signal arm of the source gets rid of the problem. Using an OI in the signal arm, however, is not a choice taken lightheartedly, since it negatively affects the transmission, lowering the heralding efficiency. Consequently, this measure results in only a small improvement on the final signal to noise ratio. Both measures described in this section, switching the source and using the OI, lead to an overall factor 2 improvement in the SNR. Even with these additional optical elements in the source setup we reach heralding efficiencies of $\eta_\text{h} = 40(4)\,\%$ including the losses due to the delay line.

In an effort to reduce the delay between the generation of the photon pair and the detection of the idler we compacted the setup. For the interfaced measurements the spectral filters and the herald SPAD were all placed in one fiber-coupled "blackout-box". A further OI in the idler arm was necessary to avoid an external cavity from forming between the plane surface of the ppKTP crystal and the idler etalon. This unwanted cavity caused some of the idler photons to be delayed by an integer multiple of its round-trip of 30 ns, triggering the experiment logic at a wrong time. Since our measurements are triggered by the detection of the idler, such a delay would appear as a 'precursor' pulse in the histogram for the arrival-time difference between signal and idler is shortened by this effect.

With these final improvements, described above, we performed the interfacing experiments with the Zeeman memory as they are described in [29] and briefly summarized

in Chapter 6. Placing an OI in the signal arm required us to make minor adjustments regarding the seed laser path. The seed laser can't be coupled in backwards through the signal fiber anymore. A flip mirror, before the OI, is added for coupling in the seed laser to align the source. By using different fibers, the seed and signal mode do not intrinsically overlap. The new seed path can be aligned by injecting an auxiliary laser beam backwards through the idler fiber (before installing the OI) and then using a quarter-wave plate to direct the backreflection from the crystal towards the fiber coupler. Furthermore, it is convenient to add a second flip mirror after the OI allowing for a 'source bypass'. This is a practical way to swap between WCPs and heralded single photons as input for the memory. This way the source and the memory can be decoupled without unplugging the 60 m fiber, even after the two setups have been interfaced. This allows them to be initialized in parallel with the seed and signal laser respectively.

3.4 Operation at the ^{87}Rb D_2 Line

This SPDC source was originally designed to emit heralded photons at the ^{87}Rb D_1 line. A closer look at the quasi-phase-matching conditions revealed that, at least theoretically, it should be possible to tune the source to the Rb D_2 line, while remaining within the tuning range of our pump laser and keeping a crystal temperature close to the specified working range. Furthermore, the dielectric coatings applied to the crystal are expected to perform similarly at 780 nm, inducing only minor changes in the cavity parameters.

In order to systematically tune the source to 780 nm we first compared the agreement of experimental parameters with the ones computed from the phase-matching conditions at 795 nm. GB noticed a discrepancy of about 0.25 nm in the pump wavelength between the predicted value and the one documented in [182]. The poling period is the only quantity for which the manufacturers did not specify an error. By slightly varying the poling period in the model we find that for $\Lambda_{pp} = 10.1465\,\mu\text{m}$ (corresponding to a deviation of $< 0.5\,\%$ from the specified value) it achieves good agreement with the experiment. Figure 3.1 already uses this adjusted value for the poling period. By computing the required pump frequency for a given temperature with this modified poling period, it was straightforward to systematically generate a DFG signal with a seed laser at the ^{87}Rb D_2 line. The prediction fits excellently within the accuracy of the spectrometer used to determine the pump wavelength. For tuning from the D_1 to the D_2 line or vice versa it is necessary to tune the angle of the pump's ECDL grating. For the operation at the D_2 line the ppKTP temperature is set around 32 °C while the pump wavelength is approximately tuned to 402.9 nm, resulting in an idler wavelength of about 833.2 nm.

By decreasing our signal wavelength to 780 nm, energy conservation dictates that the idler wavelength has to increase. The resulting new wavelength, however, lies outside of the tuning range of the IF we used so far in the idler filtering stage. In the absence of a readily available suitable filter, we tried working without it. To make sure that

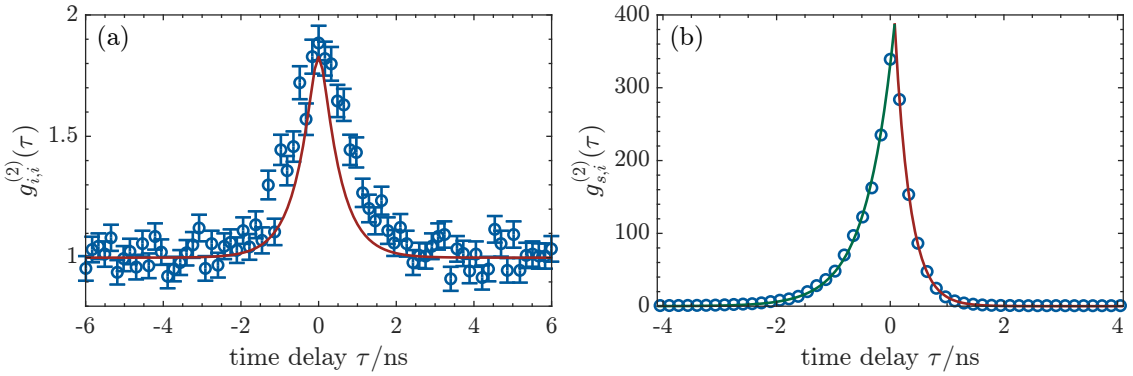

Figure 3.14: Second-order correlations of the source operating at the Rb D_2 line. **(a)** Time-resolved unconditioned idler autocorrelation with a bin-width of 162 ps. **(b)** Time-resolved, unfiltered cross-correlation for 162 ns bins.

the idler is still filtered down to a single mode we estimate the effective mode number by recording an (unconditioned) idler-idler autocorrelation. A $g^{(2)}_{i,i}(0) = 1.262(9)$ for one single 6.48 ns wide bin was recorded. Figure 3.14(a) shows the measured time-resolved autocorrelation. The autocorrelation is modeled according to Eq. (3.6), where τ_0 was obtained from the width of the cavity spectrum. These more recent measurements were performed with our new SNSPDs (see Section 5.2.8). Due to the better timing resolution of this type of detector, a peak value closer to two is reached, compared to the older measurements. The dominant jitter contribution can be attributed to the time-tagger with its time resolution of 162 ps. The reached peak value shows that the etalon alone is capable of filtering the idler to a single mode. The overall performance of the source at the D_2 line is comparable to what we measured previously on the D_1 line. For a typical pump power of 4.5 mW we record a rate of 8.44×10^5 cts/s for the unfiltered signal while the idler rate is 1.7×10^5 cts/s. The conditioned signal autocorrelation for one single 6.48 ns bin reaches a value of $g^{(2)}_{i:s,s}(0) = 0.0298(1)$. The cross-correlation shows prominent super-thermal bunching with a $g^{(2)}_{s,i}(0) = 361$ at zero time delay for a bin width of 162 ps. The asymmetry between the two exponential edges is more pronounced, compared to the operation at the D_1 line, which is unexpected. In order to exclude that this effect was induced by the etalon used to filter the idler to a single mode, we recorded an unfiltered cross-correlation, as can be seen in Fig. 3.14(b). A fit of the two exponential edges, without any need to correct for the IRF due to the small jitter, confirms the large asymmetry between signal and idler cavity damping rates, yielding 592(6) MHz and 319(8) MHz respectively. The shape of the cross-correlation can actually be used as well to determine whether the photon-pair source is running single mode [216, 109]. For this test a temporal resolution in the order of the round-trip time of the photons within the crystal is necessary. While for a multimode SPDC source the cross-correlation is expected to have a structure with a periodicity depending on the number of the round trips each photon of one pair had in the cavity, for a single-mode source a bilateral exponential decay is expected. A detailed description of the expectation of the cross-correlation's shape, also depending on degenerate or non-degenerate operation of the source, can be found in [108].

A cavity ringdown measurement was performed with the seed laser, pulsed with one of the amplitude modulators used for the memories, at the signal wavelength. This measurement confirmed the unexpectedly high signal cavity damping rate and is compatible with the results from the fit. The dielectric coatings, as the group indices of the crystal as well, are not different enough for 795 nm and 780 nm to explain this behavior. While further characterizations are necessary to understand the origin of this behavior, the source exhibits the necessary signatures of a heralded single-photon source in its second order correlations.

While operating the source at the D_2 line we noticed an unexpected behavior. The Pockels cell did not fully suppress the source emission when rotating the pump polarization. Rather, it emitted more noise. This behavior could be easily reproduced using a manually rotated half-wave plate instead of the Pockels cell. No rotation angle would fully suppress the detected counts on the signal detectors. Our interpretation is that in this configuration, other, competing, non-linear processes might be better phase-matched than in the designed operation range of the periodically poled crystal, leading to this additional emission. However, by adding a PBS in between the Pockels cell and the OPO we are able to fully switch the pump laser and therefore suppress the downconversion process. By doing so, the pump laser is not continuously incident on the crystal. In order to rule out that the resulting change in thermal load could lead to oscillations in the emission frequency, we used a fast oscilloscope to monitor whether the seed cavity modes would fluctuate. We did not observe any temporal change in the resonances frequency. Periodically turning off the pump light, however, caused a constant offset of the resonance frequency with respect to scenario in which the pump beam constantly shone on the crystal. This offset can be compensated by fine-tuning the crystal temperature. This requires some trial and error, since the exact offset depends on the trigger rate of the Pockels cell which in our case is determined by the heralding rate, which is not accessible when setting up the pump laser.

4 Atomic Vapor in the Hyperfine Paschen-Back Regime

With the goal in mind to isolate a nearly ideal three-level system in a hot atomic vapor, I explored the hyperfine Paschen-Back regime with spectroscopy experiments. This chapter will start with a brief theoretical description of ^{87}Rb in a strong external magnetic field. After familiarizing with this new regime by performing some simple spectroscopic measurements, I started laying the ground work for the quantum memory by investigating EIT in these conditions. The experiments were directly performed in a micro-fabricated vapor cell I received from our colleagues from Gaetano Mileti's research group in Neuchâtel. At the end of the chapter I will address the challenge of preparing the initial atomic state in a high magnetic field and present the first steps I performed towards optically polarizing the nuclear spin of the whole ensemble. The experimental results presented in this chapter are also published in [66].

4.1 Atomic Hamiltonian in a DC magnetic field

The Hamiltonian for an atom in an external static magnetic field is given by

$$\hat{H} = \hat{H}_0 + \hat{H}_{\text{hfs}} + \hat{H}_Z \tag{4.1}$$

where $\hat{H}_0$ describes the coarse atomic structure including the fine structure, $\hat{H}_{\text{hfs}}$ describes the hyperfine structure interaction, and $\hat{H}_Z$ describes the atomic interaction with the external magnetic field. The hyperfine interaction Hamiltonian is

$$\hat{H}_{\text{hfs}} = A_{\text{hfs}}\, \hat{\mathbf{I}} \cdot \hat{\mathbf{J}} + \hat{H}_{qp} \tag{4.2}$$

with the nuclear spin operator $\hat{\mathbf{I}}$ and the total electronic angular momentum operator $\hat{\mathbf{J}}$. A_{hfs} is the magnetic dipole constant, also referred to as hyperfine coupling-coefficient, and $\hat{H}_{qp}$ is the electric quadrupole Hamiltonian, which describes the electric quadrupole interaction between the nucleus and the electron[1]. This latter term vanishes for $J = 0, \frac{1}{2}$, since the electron charge distribution is spherically

[1] Technically, there is also a term describing the magnetic octupole in $\hat{H}_{\text{hfs}}$ for the excited states of the D_2 manifold. The effect, however, is so small and the measurements performed so far were not precise enough to obtain nonzero values for the octupole constant C_{hfs} [217].

symmetric in these cases [218]. The same is valid for $I = 0, \frac{1}{2}$. In general, the electric quadrupole Hamiltonian has the form [218]

$$\hat{H}_{qp} = B_{\text{hfs}} \frac{3\left(\hat{\mathbf{I}}\cdot\hat{\mathbf{J}}\right)^2 + \frac{3}{2}\left(\hat{\mathbf{I}}\cdot\hat{\mathbf{J}}\right) - I\,(I+1)\,J\,(J+1)}{2I\,(2I-1)\,J\,(2J-1)} \tag{4.3}$$

where B_{hfs} is the electric quadrupole constant. The energy shift induced by the hyperfine interaction is [218]

$$\Delta E_{\text{hfs}} = \frac{A_{\text{hfs}}}{2}K + \frac{B_{\text{hfs}}}{4}\frac{\frac{3}{2}K(K+1) - 2I(I+1)J(J+1)}{I(2I-1)J(2J-1)} \tag{4.4}$$

with $K = F(F+1) - I(I+1) - J(J+1)$.

The last term in Eq. (4.1) is the one describing the interaction of an atom with the magnetic field. This Hamiltonian has a linear and quadratic term in $\mathbf{B}$ [218]. For states with low principal quantum number n in typical magnetic fields generated in laboratories, the induced quadratic shifts are small and therefore difficult to measure. Usually, the quadratic interaction can be neglected. With higher n states, e.g. when dealing with Rydberg atoms, the quadratic effect becomes evident [219]. For a state with $n \approx 36$ in a magnetic field of 0.6 T the quadratic interaction begins to be significant [220, 219]. Consequently, for our work we can safely consider only the term linear in $\mathbf{B}$ of the magnetic interaction Hamiltonian. If the energy shift induced by the magnetic field is small compared to the fine-structure splitting, J is a 'good'[2] quantum number and the magnetic interaction Hamiltonian can be written as

$$\hat{H}_Z = \mu_B \left(g_J \hat{\mathbf{J}} + g_I \hat{\mathbf{I}}\right)\cdot\mathbf{B}\,. \tag{4.5}$$

Here g_J and g_I are the g-factors for the total angular momentum of the electron and for the nucleus respectively. The Landé factor g_J for the various atomic states can be computed by [218]

$$g_J = g_L \frac{J(J+1) - S(S+1) + L(L+1)}{2J(J+1)} + g_S \frac{J(J+1) + S(S+1) - L(L+1)}{2J(J+1)}\,. \tag{4.6}$$

The necessary constants for solving the Hamiltonians above, as well as the computed values of g_J for the states of interests for this thesis, are listed in Tab. A.1.

Different Magnetic Regimes We distinguish between different magnetic regimes based on the relative strength of the various interactions. In the hyperfine linear Zeeman (HLZ) regime the magnetic interaction can be treated as a perturbation of the hyperfine interaction. The energy splitting induced by the external magnetic

[2] In the strict sense, the eigenvalues of an operator are considered to be a good quantum number when the associated observable is a constant of motion, i.e. it commutes with the Hamiltonian. As long as the deviation from this case is small, the quantum number can be considered to be good enough.

field ΔE_Z is smaller than the hyperfine splitting ΔE_{hfs}. In this weak-field regime the nuclear spin $\mathbf{I}$ and the total electronic angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$ are coupled, resulting in the total angular momentum $\mathbf{F} = \mathbf{I} + \mathbf{J}$. In this regime the Hamiltonian eigenstates coincide, to a good degree, with the total angular momentum eigenstates $|F, m_F\rangle$. The hyperfine levels can take the values $|J-I| \leq F \leq J+I$. In this regime, each hyperfine level F is split symmetrically around the zero field level into $2F+1$ different m_F-levels, where m_F is the projection of $\mathbf{F}$ along the direction of the applied magnetic field. The Zeeman effect shifts the levels by $\Delta E_Z = g_F m_F \mu_B B$, assuming a magnetic field along the z-direction.

For higher fields, where the magnetic interaction becomes larger than the hyperfine interaction, but remains smaller than the fine-structure interaction, we enter the HPB regime. The nuclear spin and total angular momentum of the electron decouple. As a consequence, F and m_F are not good quantum numbers anymore. Instead, it is convenient to describe the system in terms of J, m_J, I, and m_I. Since J and I are constant within one fine-structure level, I will use the compressed state notation $|m_J, m_I\rangle = |J, m_J, I, m_I\rangle$. For each value of m_J there are $2I+1$ m_I-levels. The spacing of these levels is proportional to m_J and m_I. In fact, in this regime the energy shift induced by the external field is described by $\Delta E_Z = (g_J m_J + g_I m_I)\mu_B B$. For intermediate fields there are no good quantum numbers to describe the interaction.

If the magnetic field is increased ever further, the energy shifts reach the order of the fine-structure splitting and beyond. This leads to the electron spin $\mathbf{S}$ and the orbital angular momentum of the electron $\mathbf{L}$ to decouple as well, rendering J a bad quantum number to describe the system. However, for Rb, magnetic fields $> 218\,\text{T}$ are required to reach this so called fine Paschen-Back (FPB) regime [220] – far above what we can generate in our laboratory. In this scenario Eq. (4.5) needs be replaced with the fully uncoupled one

$$\hat{H}_Z = \mu_B \left(g_L \hat{\mathbf{L}} + g_S \hat{\mathbf{S}} + g_I \hat{\mathbf{I}} \right) \cdot \mathbf{B} \,. \tag{4.7}$$

Furthermore, the term in $\hat{H}_0$ describing the fine-structure interaction needs to be considered, since it is proportional to $\hat{\mathbf{L}} \cdot \hat{\mathbf{S}}$.

The necessary magnetic field strength to enter the HPB regime depends on the hyperfine interaction of the specific investigated atomic species. In general, the condition $B \gg B_0 = A_{\text{hfs}}^{GS}/\mu_B$ [218] can be obtained by comparing the energy shifts induced by the hyperfine and the magnetic interaction. It is noteworthy that for the the ground state of alkali atoms the following relation holds $A_{\text{hfs}}^{GS} = h\nu_{\text{hfs}}/(I+1/2)$ [218], where ν_{hfs} is the ground-state splitting. For ^{133}Cs [221, 222], ^{87}Rb and ^{85}Rb [220, 223] $B_0 \approx 170\,\text{mT}, 240\,\text{mT}$ and $72\,\text{mT}$ respectively. In the case of ^{39}K $B_0 \approx 17\,\text{mT}$, meaning that the HPB regime is reached with magnetic fields weaker by one order of magnitude compared to ^{87}Rb [224]. For the excited states the effect of the applied magnetic field begins even at lower fields. In fact, for ^{7}Li the FPB regime for the D lines was investigated experimentally with fields of $< 1\,\text{T}$, observing the mixing of different fine-structure levels [225].

In the following two chapters I will focus on the HPB regime. Consequently, when talking about coupled/uncoupled bases I will be referring to the coupling of **I** and **J**.

4.2 ^{87}Rb D Lines

In this section I will concentrate on the states involved in the Rb D line transitions and their behavior in an external magnetic field. Rubidium naturally occurs in the form of the two isotopes ^{85}Rb (72.2 %) and ^{87}Rb (27.8 %) [226], which have a nuclear spin of $I = \frac{5}{2}$ and $I = \frac{3}{2}$ respectively. In our experiments we will work with ^{87}Rb due to its smaller number of Zeeman sublevels and its larger ground-state splitting, especially relevant when working without applied magnetic field.

The transitions between the ground states and first excited states of alkali metal atoms are referred to as the D lines. Due to the fine-structure splitting, the lowest excited state is split in two, giving origin to two lines: 1) the D_1 line describing the transition $n\,^2S_{1/2} \rightarrow n\,^2P_{1/2}$ and 2) the D_2 line describing the transition $n\,^2S_{1/2} \rightarrow n\,^2P_{3/2}$, where n is the principal quantum number of the valence electron. For Rb the principal quantum number has a value of $n = 5$. Note that the name of the D lines is of historic origin and does not imply the involvement of the higher orbital D-states.

The zero detuning frequency for ^{87}Rb, without considering the hyperfine splitting, is roughly 377.11 THz and 384.23 THz [217] for the D_1 and D_2 lines respectively. The fine-structure splitting is large enough that the two D line components usually are treated separately. Commonly, in fact, different lasers are needed experimentally to address the one or the other line as they are separated by ~ 15 nm. This difference is much larger than the energy splittings the magnetic field induces in the HPB regime. We can thus neglect the first term from the Hamiltonian in Eq. (4.1) for the purpose of the following considerations. This term would result in the constant offset given by the zero detuning frequency mentioned above.

For the ^{87}Rb $5\,^2S_{1/2}$ ground state the atomic Hamiltonian for an external field along the z-direction $\mathbf{B} = \mathbf{e}_z B_z$, where $\mathbf{e}_z$ is the unit vector along the z-axis, thus simplifies to

$$\hat{H} = A_{\mathrm{hfs}}\,\hat{\mathbf{I}}\cdot\hat{\mathbf{J}} + \mu_B B_z\left(g_J\hat{J}_z + g_I\hat{I}_z\right) \tag{4.8}$$

with $J = \frac{1}{2}$ and $I = \frac{3}{2}$. The $5\,^2P_{1/2}$ state is described by the same Hamiltonian, with merely the constants changing. For the $5\,^2S_{1/2}$ and $5\,^2P_{1/2}$ states the quadrupole term is zero. For the Hamiltonians and a table listing the important constants for the $5\,^2P_{3/2}$ term I direct the reader to Appendix A. For an even more extensive collection of Rb data I refer the reader to Daniel Steck's alkali data [217], which are also available for other alkali metals.

By diagonalizing the Hamiltonian we obtain the eigenvalues, which correspond to the energies of the respective states. For terms with $J = \frac{1}{2}$, i.e. $5\,^2S_{1/2}$ and $5\,^2P_{1/2}$,

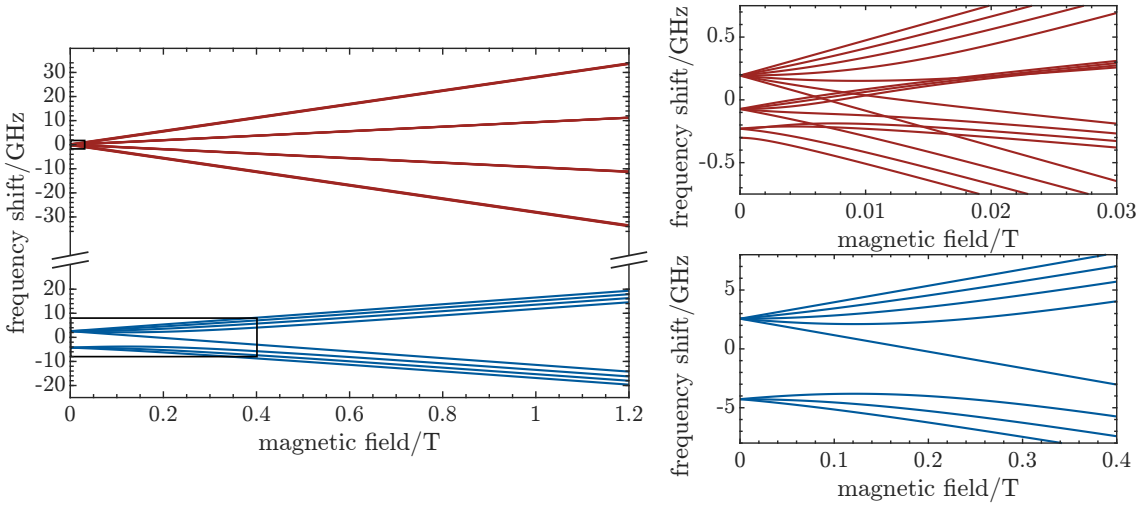

Figure 4.1: Breit-Rabi diagrams for the $5\,^2S_{1/2}$ (blue) and $5\,^2P_{3/2}$ (red) states of ^{87}Rb. Frequency shift is relative to the zero detuning frequency of the respective fine-structure level. The regions marked by a black box are shown magnified in the insets on the right-hand side. They show the crossing of the energy levels in the weak field regime.

the Hamiltonian can be diagonalized analytically by using the Breit-Rabi formula [227] or, as I did, numerically, which is anyhow the only viable approach for the $5\,^2P_{3/2}$ excited state. In order to obtain the trend of the energies as a function of the field, the eigenvalues need to be computed for each value of B. Figure 4.1 shows the Breit-Rabi diagram, which is a representation of the energy shift of the levels as a function of the applied external field, for the terms making up the D_2 line. The panels on the right-hand side show magnifications of the magnetic field regions where the atomic levels rearrange energetically. Note that for the excited state this happens at fields an order of magnitude smaller than for the ground state. The resulting energy level diagram at large fields is schematically shown in Fig. 4.2. For an external magnetic field of 1.06 T the ground-state splitting ranges from about 25 GHz to 35 GHz for $m_I = -\frac{3}{2}$ and $m_I = +\frac{3}{2}$ respectively (compare to the 6.8 GHz splitting at zero field). The splittings between adjacent m_J-manifolds of the $5\,^2P_{3/2}$ term are approximately 20 GHz, while the splitting between the different m_I levels within one m_J-manifold are in the order of single GHz and tens to hundreds of MHz for the ground and $P_{3/2}$-excited state respectively. An energy diagram with all the frequency splittings for the D_2 line can be found in Appendix A.

In general, each hyperfine state is a superposition of different $|m_J, m_I\rangle$ sublevels with the same total angular momentum projection $m_F = m_J + m_I$. For the ground state $5\,^2S_{1/2}$ the general superpositions of $|m_J, m_I\rangle$ states composing the various energy levels are listed in Tab. 4.1. In Fig. 4.2 gray lines show the states that are coupled by hyperfine interaction. The generalized Clebsch-Gordon coefficients a_i and b_i can be determined from the diagonalized Hamiltonian. They correspond to the non-zero entries of the eigenvectors. The stretched states, $|F = 2, m_F = \pm 2\rangle \leftrightarrow |m_J = \pm\frac{1}{2}, m_I = \pm\frac{3}{2}\rangle$, consist of a single $|m_J, m_I\rangle$ state for any given field value.

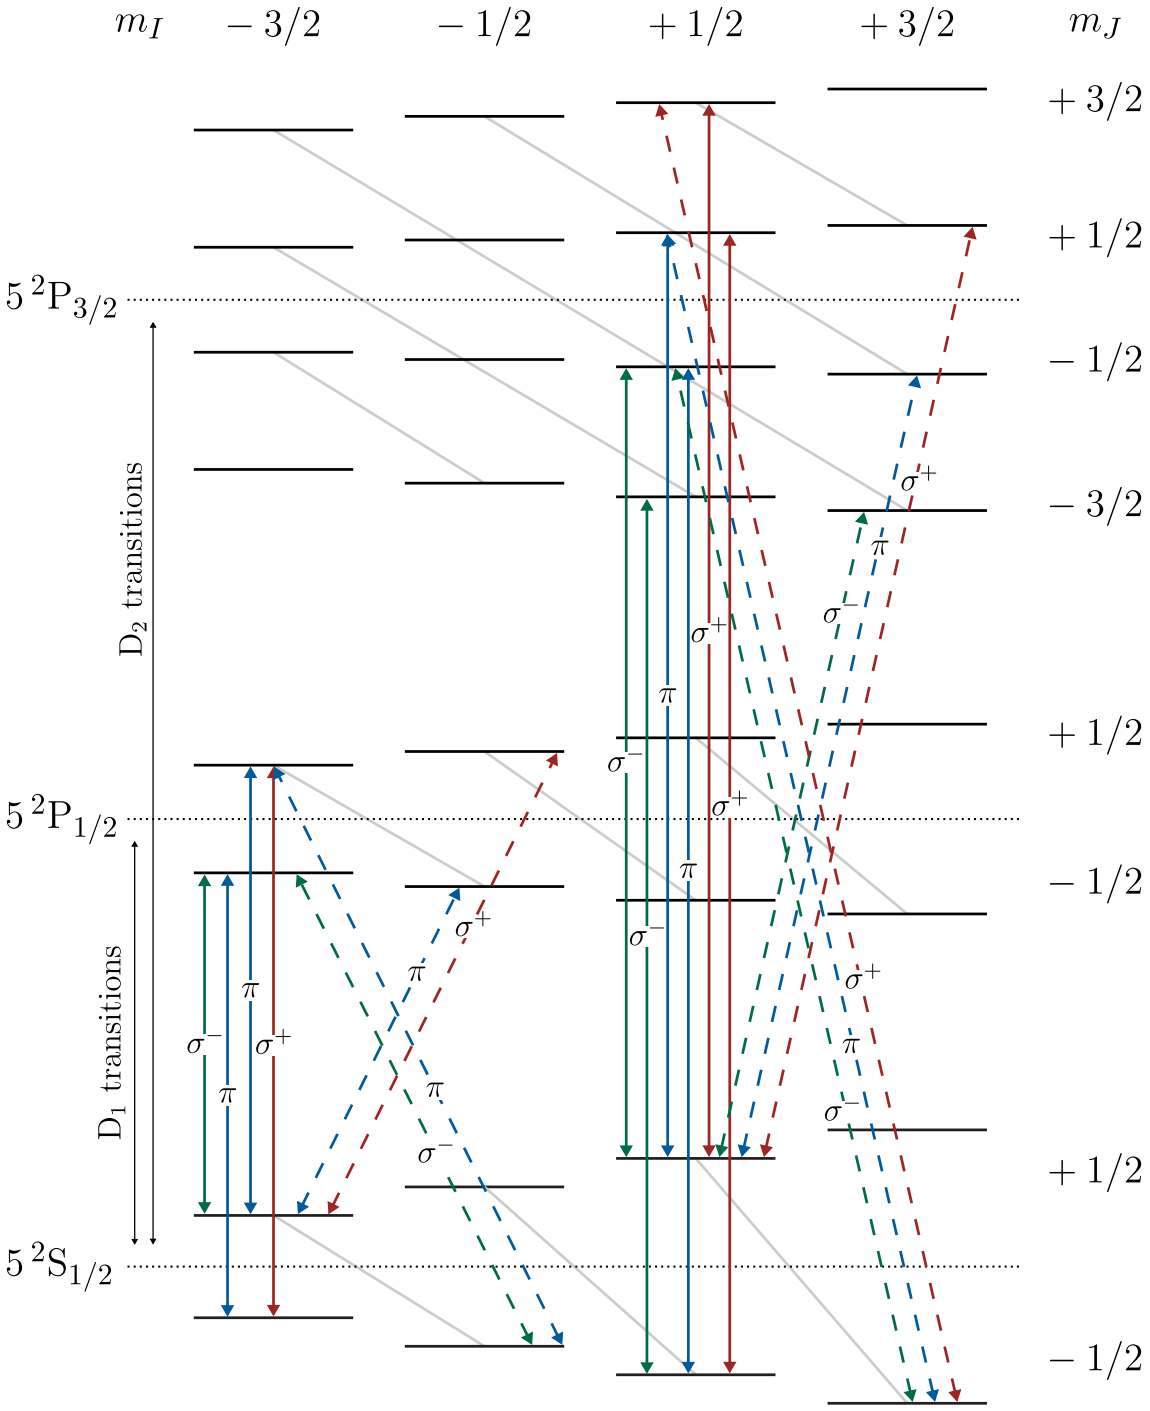

Figure 4.2: Energy levels of ^{87}Rb in the HPB regime represented in the $|m_J, m_I\rangle$ basis. The coupling between the sublevels is indicated by the gray lines. Selected transitions are shown. The allowed D_1 and D_2 transitions for one m_I manifold each, $-\frac{3}{2}$ and $+\frac{1}{2}$ respectively, are represented as solid lines. σ^+, π, and σ^- transitions are color coded (red, blue, green). Some singly forbidden transitions arising from the residual coupling of the ground state are displayed as dashed lines. The transitions for the ground-state combinations corresponding to $m_F = -1$ and $m_F = 1$ (in weak fields) are shown for the D_1 and D_2 respectively. Singly forbidden transitions arising from the residual coupling of the $5\,^2\mathrm{P}_{1/2}$ term and doubly forbidden transitions are not shown. The energy splittings are not drawn to scale.

Level	Weak field $\lvert F, m_F\rangle$	Intermediate regime $\lvert m_J, m_I\rangle$	Strong field $\lvert m_J, m_I\rangle$
A_1	$\lvert 1,+1\rangle$	$a_1 \lvert -\frac{1}{2},+\frac{3}{2}\rangle + b_1 \lvert +\frac{1}{2},+\frac{1}{2}\rangle$	$\lvert -\frac{1}{2},+\frac{3}{2}\rangle$
A_2	$\lvert 1,0\rangle$	$a_2 \lvert -\frac{1}{2},+\frac{1}{2}\rangle + b_2 \lvert +\frac{1}{2},-\frac{1}{2}\rangle$	$\lvert -\frac{1}{2},+\frac{1}{2}\rangle$
A_3	$\lvert 1,-1\rangle$	$a_3 \lvert -\frac{1}{2},-\frac{1}{2}\rangle + b_3 \lvert +\frac{1}{2},-\frac{3}{2}\rangle$	$\lvert -\frac{1}{2},-\frac{1}{2}\rangle$
A_4	$\lvert 2,-2\rangle$	$\lvert -\frac{1}{2},-\frac{3}{2}\rangle$	$\lvert -\frac{1}{2},-\frac{3}{2}\rangle$
A_5	$\lvert 2,-1\rangle$	$a_5 \lvert +\frac{1}{2},-\frac{3}{2}\rangle + b_5 \lvert -\frac{1}{2},-\frac{1}{2}\rangle$	$\lvert +\frac{1}{2},-\frac{3}{2}\rangle$
A_6	$\lvert 2,0\rangle$	$a_6 \lvert +\frac{1}{2},-\frac{1}{2}\rangle + b_6 \lvert -\frac{1}{2},+\frac{1}{2}\rangle$	$\lvert +\frac{1}{2},-\frac{1}{2}\rangle$
A_7	$\lvert 2,+1\rangle$	$a_7 \lvert +\frac{1}{2},+\frac{1}{2}\rangle + b_7 \lvert -\frac{1}{2},+\frac{3}{2}\rangle$	$\lvert +\frac{1}{2},+\frac{1}{2}\rangle$
A_8	$\lvert 2,+2\rangle$	$\lvert +\frac{1}{2},+\frac{3}{2}\rangle$	$\lvert +\frac{1}{2},+\frac{3}{2}\rangle$

Table 4.1: Summary of the general combinations of $|m_J, m_I\rangle$ states composing the various levels of the $5\,^2S_{1/2}$ term. The levels are labeled $A_1 \to A_8$ in order of increasing energy.

Since they have constant magnetic moments, for they have a fixed combination of m_I and m_J, they appear as a straight line in the Breit-Rabi diagrams. In the absence of an external magnetic field the coefficients a_i and b_i correspond to the Clebsch-Gordon coefficients. For an increasing magnetic field the coefficients follow the trend $a_i \to 1$ and $b_i \to 0$.

Figure 4.3 shows the changing composition of the hyperfine eigenstates of the $5\,^2S_{1/2}$ term with increasing applied magnetic field. The states are ordered as in the case of the weak field regime $|F, m_F\rangle$, where the lower row corresponds to values of $F = 1$ and the upper one to $F = 2$. Within one column of the plot each color represents the coefficients for one specific sublevel. The corresponding table of state combinations and plots of the magnetic field dependent generalized coefficients for the excited states can be found in the Appendix A.

For the ground state, even with a field as high as 1.2 T the mixture of the sublevels is still not negligible, with the b_i coefficients having still a value of about 0.1. This implies that the nuclear spin and the total angular momentum of the electron are not fully decoupled under those conditions. This residual coupling leads to 'forbidden' transitions that will be discussed in the next section and are of experimental interest, as explained in Section 4.5.3. The $5\,^2P_{1/2}$ term has coefficients about a factor 3 smaller than those of the ground state (cf. Fig. A.2), while for the $5\,^2P_{3/2}$ exited state all the coefficients but the dominant one can be neglected to a good approximation at the considered fields (cf. Fig. A.3). This difference is a direct consequence of the P terms having smaller magnetic dipole constants than the ground state.

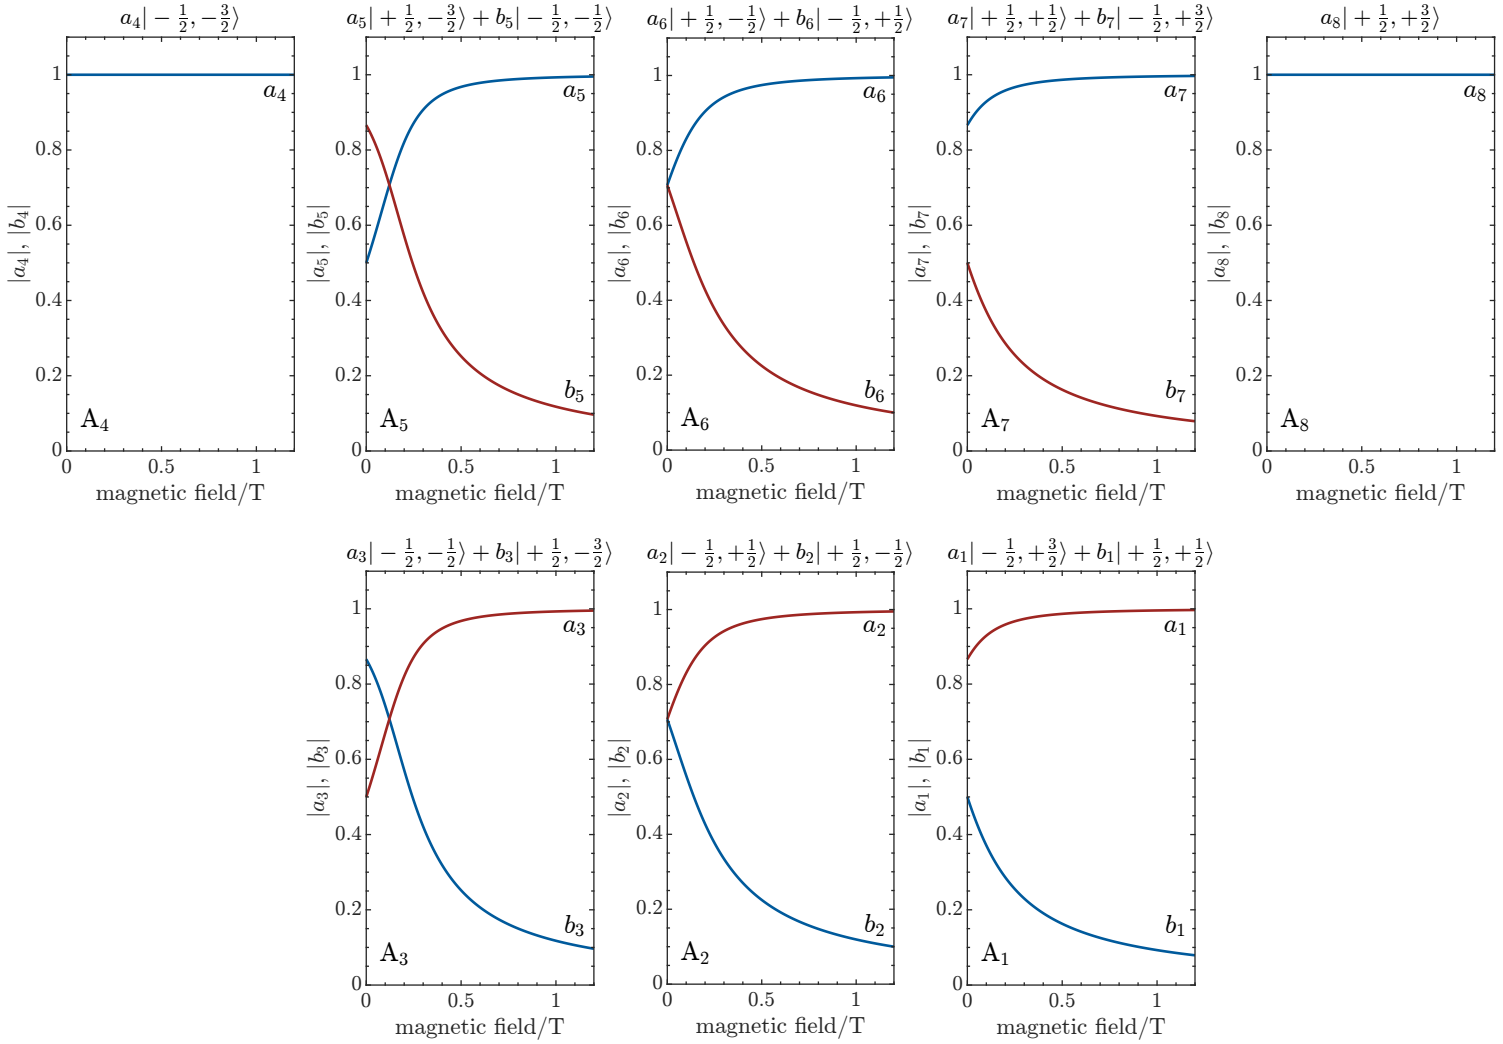

Figure 4.3: Coupling coefficients for the $5\,^2\mathrm{S}_{1/2}$ hyperfine levels in the uncoupled $|m_J, m_I\rangle$ basis as a function of the applied magnetic field. The sublevels are arranged as they usually are in the weak field regime, where the bottom row corresponds to $F = 1$ and the top one to $F = 2$. Within one column each color represents one specific sublevel.

4.3 Atom-Light Interaction

In the absence of the hyperfine coupling, optical dipole transitions couple only to the total angular momentum of the electron. These transitions, which I will refer to as 'allowed', obey the selection rules $\Delta m_I = 0$ and $\Delta m_J = 0, \pm 1$, for π and $\sigma^{\pm}$ polarized light respectively. These transitions occur as quadruplets (one for every value of m_I). In the energy scheme depicted in Fig. 4.2 the allowed transitions appear as (solid) vertical lines. The allowed transitions for the D_1 and D_2 lines are exemplified for the cases of $m_I = -\frac{3}{2}$ and $m_I = +\frac{1}{2}$ respectively. They follow similarly for the other manifolds.

The calculations of the ground-state coefficients show that, even at the levels of the magnetic field we are working with, the hyperfine coupling is not negligible. For the ground state levels, I and J are not fully decoupled; for the D_1 excited states there is a small coupling left (experimentally already difficult to see) while for the D_2 states coupling is negligible. Transitions with $\Delta m_I \neq 0$ can be thus excited through an indirect interaction of the light with the nucleus through the residual hyperfine coupling. Since the atom-light system conserves the total angular momentum, these transitions obey the relations $\Delta m_I + \Delta m_J = 0, \pm 1$ for π and $\sigma^{\pm}$ polarized light, respectively. I will refer to the transitions with $|\Delta m_I| = 1$ as 'singly forbidden' transitions. In Fig. 4.2 the singly forbidden transitions for both the D_1 and the D_2 line are illustrated as dashed lines for the coupled sublevel pair with $m_I + m_J = -1$ and $m_I + m_J = +1$ respectively. For the other coupled ground-state sublevels the transitions follow analogously. These transitions appear as multiplets of three. The figure shows only singly forbidden transitions arising from the residual coupling of the ground state. Transitions with $|\Delta m_I| = 2$ ('doubly forbidden'), or $|\Delta m_I| = 3$ ('triply forbidden'), can also occur, but are suppressed even more strongly and are therefore not shown in the figure. By optical pumping along forbidden transitions, the nuclear spin of an atomic vapor can be directly polarized [221]. I will come back to this technique at the end of this chapter for it is of interest regarding the atomic state preparation in our system.

Since I will show and use simulated atomic spectra throughout this thesis, I will briefly summarize how they can be computed. A more rigorous description on how to compute a spectrum for a given set of parameters can be found in [228, 229, 230]. Under the assumption of a near-resonant, weak probe laser field, the various optical transitions can be considered separately as a two-level system. The absorption profile of an individual hyperfine transition can be represented by a single Voigt profile. The Voigt profile [228] is the convolution of a Lorentzian distribution describing the natural lineshape of the atoms and any additional homogeneous broadening with the Gaussian distribution corresponding to the inhomogeneous Doppler-broadening arising from the different atomic velocities in the hot vapor. Additional homogeneous broadening mechanisms such as pressure broadening, induced by alkali-alkali or alkali-buffer gas collisions, simply modify the width of the Lorentzian profile. Each line profile is centered on the corresponding transition frequency and is scaled by its transition strength. Summing over all transitions ultimately results in the total

spectrum. By diagonalizing the atomic Hamiltonian for the ground and the excited states, the eigenenergies and the eigenstates are obtained. The energy difference between the ground state $|g\rangle$ and excited state $|e\rangle$ involved in the transition yields the transition frequency. Generally, the eigenstates are a combination of the uncoupled basis states. The transition strengths are proportional to the electric dipole matrix elements squared $|\langle g|e\hat{\mathbf{r}}|e\rangle|^2$. Finally, by adding up all the single transitions the total electric susceptibility is obtained, where the real and imaginary parts describe the dispersion and absorption of the multi-level system respectively. For computing the spectra I used both a Matlab script written by Andrew Horsley and the open-source Python program ElecSus developed in Ifan Hughes' research group at Durham University [229, 231]. I used the latter especially when determining atomic parameters, such as temperature, by fitting measured spectra.

Following the method described above, the spectrum of the D_2 transitions for an external magnetic field of 1.06 T is computed for the atomic parameters corresponding to the vapor cell used for the experiments performed in this work. A 2 mm long vapor cell filled with enriched ^{87}Rb with an estimated abundance of 90 % at 97 °C is considered. Furthermore, additional homogeneous broadening due to 10 mbar of Ar buffer gas (derived from the experiments described in Appendix C) is accounted for. The spectrum is plotted in Fig. 4.4 for σ^- (green), π (blue) and σ^+ (red) polarized light. The arrows in the insets show the sublevels involved in the corresponding transitions. The allowed transitions are the strongest, resulting in the highest resonant ODs. In the energy-level schemes they are represented as solid lines. The energy diagrams are ordered as the transition multiplets appear in the spectrum. The residual ^{85}Rb component in the cell induces absorption at frequencies corresponding to the center of each quadruplet. The effect can be identified by the transparency between the two central transitions of each quadruplet to being lower and the transitions themselves having a slightly higher absorption with respect to the outer transitions of the same manifold. The ODs of the singly forbidden transitions originating from the hyperfine coupling in the ground state are two orders of magnitude lower than that of the allowed transitions. In the level schemes the involved atomic states are connected by dashed arrows. The transitions arising from the remaining mixing in the $5\,^2P_{3/2}$ term are weak. In the plot they can be barely identified at about ± 15 GHz on the blue curve.

Figure 4.5 shows the corresponding spectrum for the D_1 line. Here, the hyperfine coupling of the $5\,^2P_{1/2}$ term is large enough that even the residual hyperfine coupling of the excited state leads to recognizable transitions. These are weaker than the singly forbidden transitions originating from the residual coupling in the ground state. Doubly forbidden transitions, appearing as doublets, can be seen as well in the computed spectrum. These transitions originate from the remaining couplings of both the $5\,^2S_{1/2}$ and $5\,^2P_{1/2}$ terms. However, as one might expect, they are considerably weaker, with about four orders of magnitude lower ODs than the allowed transitions.

In the following, "forbidden transitions" refers to singly forbidden transitions originating from the ground state coupling, since all other forbidden transitions are too

weak to be observed in our measurements.

4.4 Experimental Setup

Having reviewed the theoretical behavior of the ^{87}Rb D lines in an external magnetic field, I continue here with some experimental measurements to lay the ground for the quantum memory implementation. To explore and familiarize with Rb in the HPB regime I start off with an absorption spectroscopy experiment. Then I add further light fields to investigate optical pumping and EIT. In the following sections I will start by describing in detail the key components of the experiment and subsequently show the whole spectroscopy setup.

4.4.1 Electromagnet

Dictated by the geometry of the utilized electromagnet, the high magnetic field is applied perpendicularly to the optical axis, parallel to the optical table. This configuration is known as Voigt geometry. Choosing the direction of the magnetic field as our quantization axis, π-polarized light translates to linearly horizontally polarized light in the lab frame, whereas σ^+ and σ^- polarization correspond to an electric field right- and left-circularly polarized in the plane perpendicular to the magnetic field respectively. However, for the considered geometry these are not valid polarizations of light propagating along the optical axis, since an electromagnetic wave is a transverse wave. Linearly vertically polarized light in the lab frame can, however, be decomposed into the equal superposition of σ^+ and σ^- polarizations. For our experiments this means that we can only apply both polarizations simultaneously. As we have seen in the previous sections, however, this should not pose a problem for addressing single transitions since the σ^+ and σ^- transitions are separated by large frequency differences. When the light is resonant to one transition then the other polarization component is more than 20 GHz off-resonant (for the allowed transitions) in a tesla-order field. This does, however, result in a reduction by a factor of $\sqrt{2}$ of the Rabi frequency. Most other works investigating hot vapors in high magnetic fields are performed in the Faraday geometry, where the magnetic field is oriented along the optical axis [220, 221, 232]. In this geometry no polarization component of the light couples to π transitions.

To generate the Tesla-order magnetic fields for the experiments in this thesis I used a BRUKER B-E 10 (made in 1988) electromagnet with a BRUKER B-MN60/20C4 (made in 1979) power supply. These devices had been decommissioned before being rediscovered in the basement of our department and put to good use for the frequency-tunable microwave imaging experiments Andrew Horsley set up for his thesis [168]. The electromagnet is composed of a pair of coils wound around cylindrical ferromagnetic cores. The cores can be equipped with tapered nose-cones with a 4.5 cm cross-section at the end.

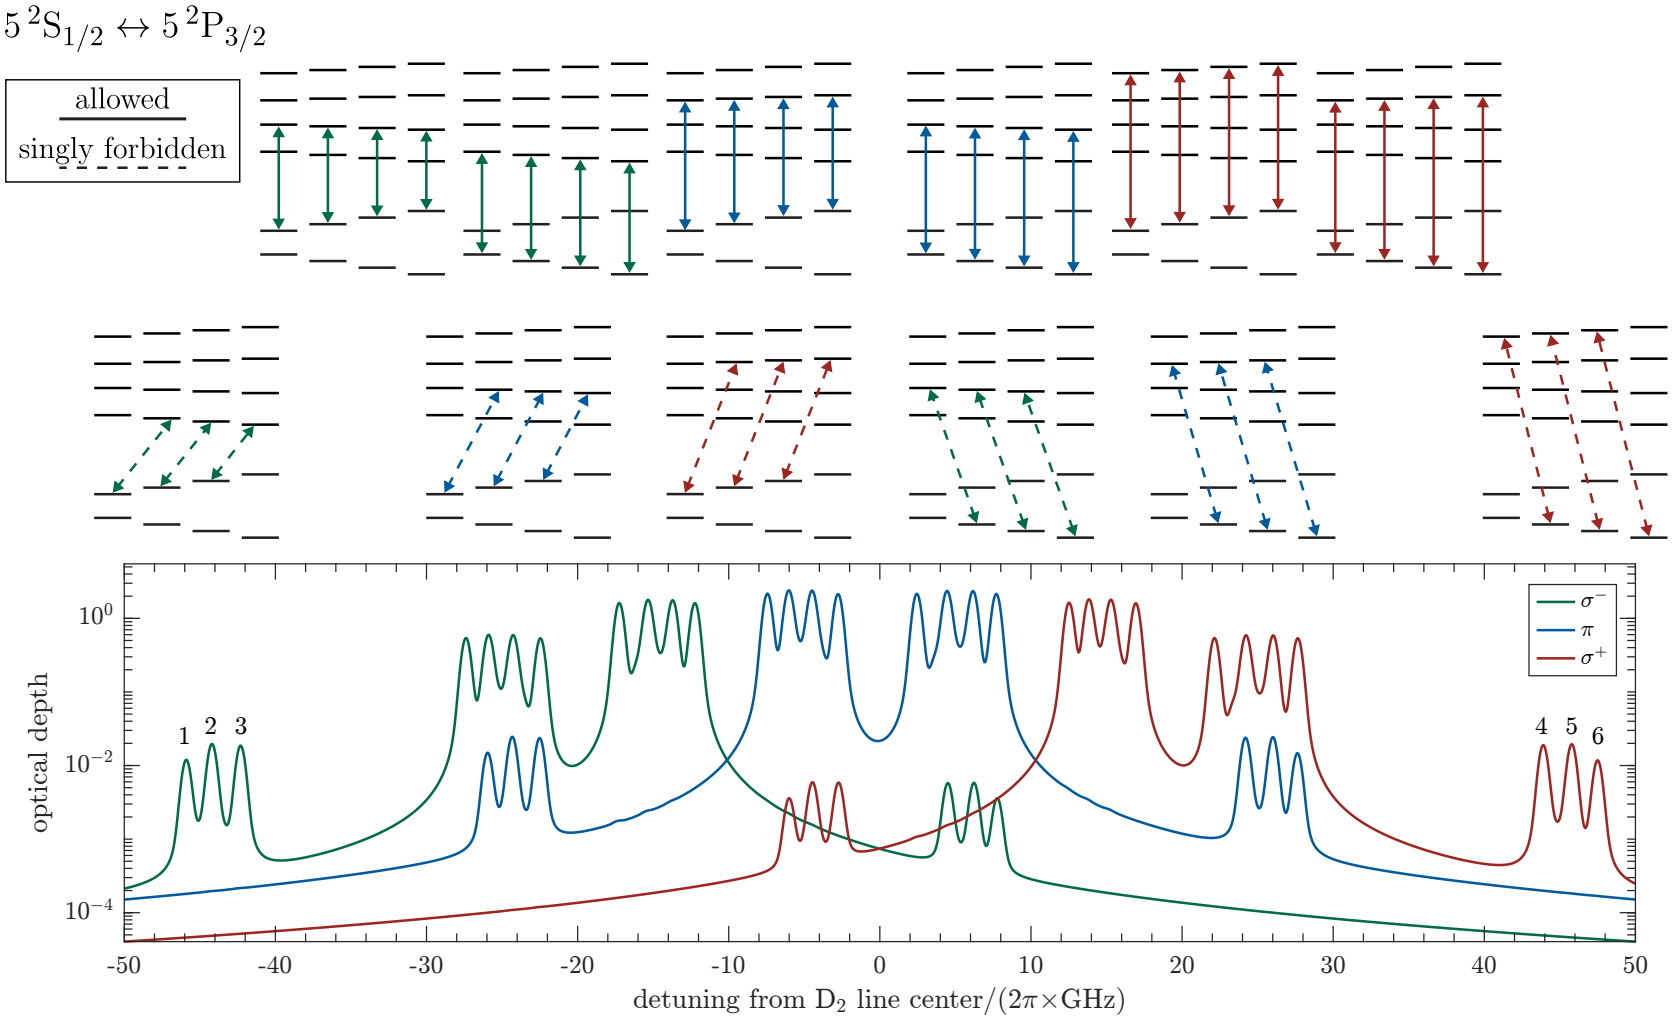

Figure 4.4: Computed spectrum of the ^{87}Rb D_2 line in an external magnetic field of 1.06 T. The considered cell is 2 mm long and contains 90 % enriched ^{87}Rb with 10 mbar of Ar. The atoms are heated to 97 °C. The color indicates the polarization. The level-scheme insets show the involved sublevels for each manifold of transitions. They are roughly ordered by frequency and OD. Allowed transitions are represented by solid arrows while singly forbidden ones by dashed arrows. The labels 1–6 are used to identify the transitions used for pumping the nuclear spin.

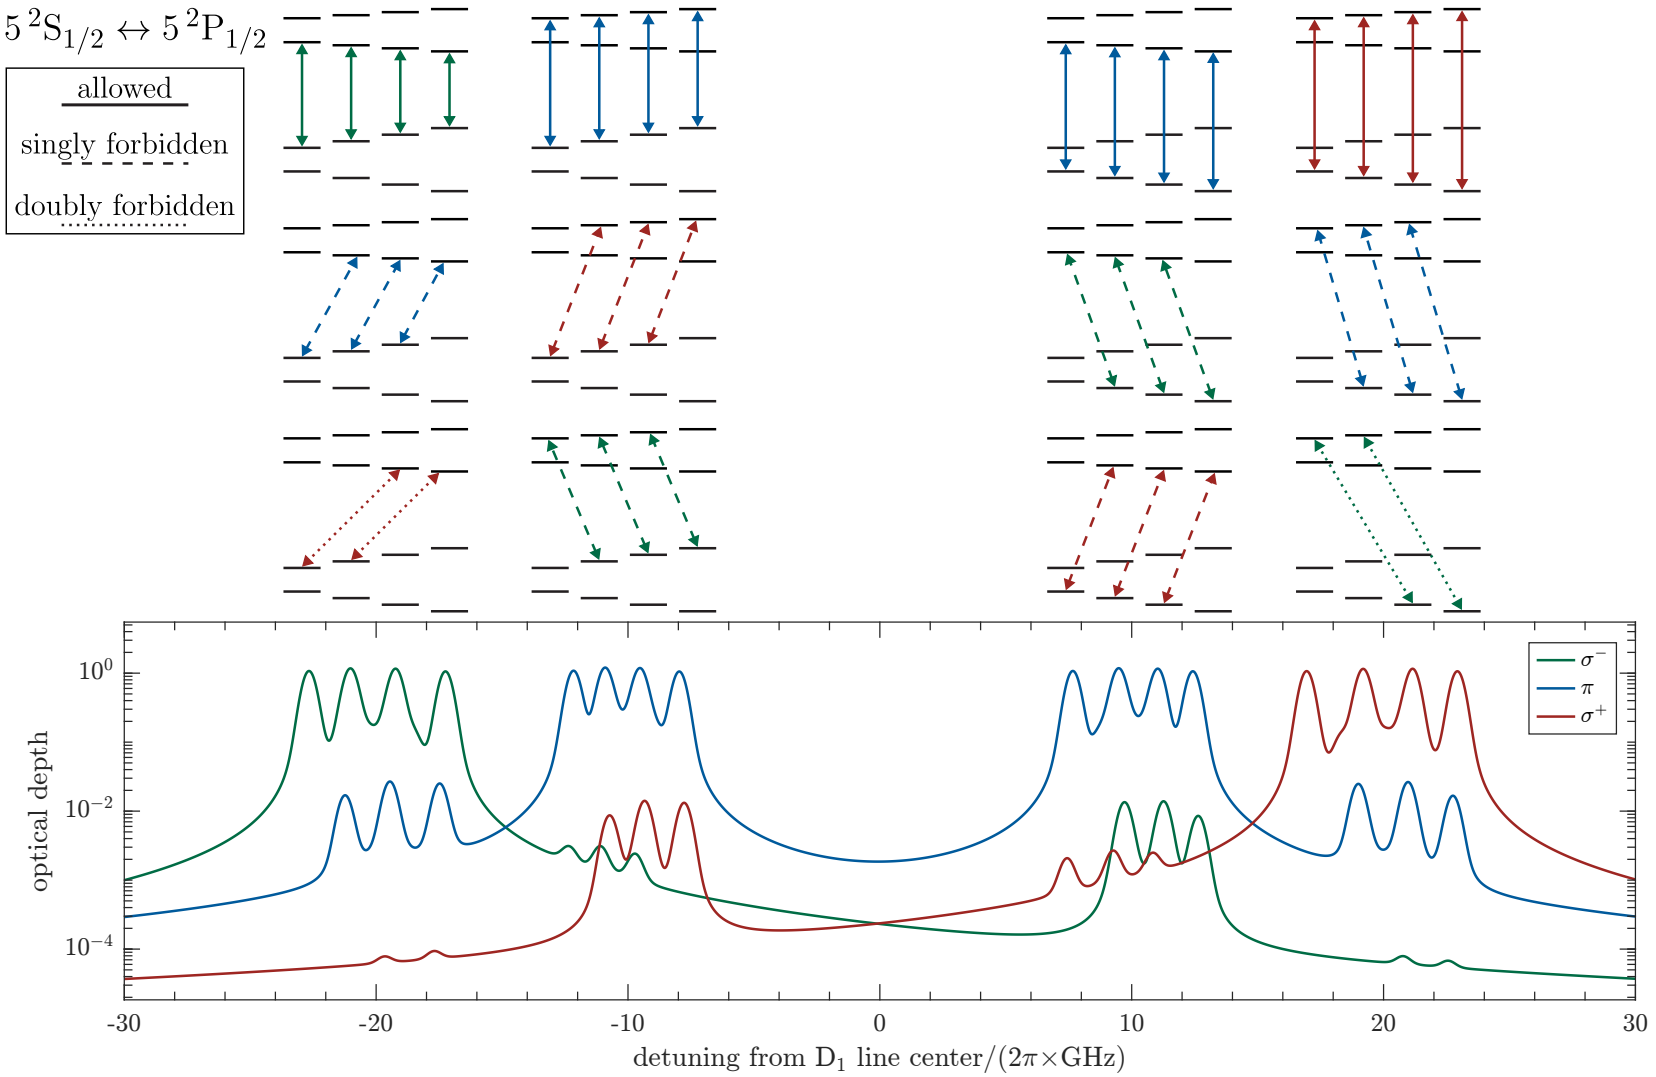

Figure 4.5: Computed spectrum of the ^{87}Rb D_1 line in an external magnetic field of 1.06 T. The considered cell is 2 mm long and contains 90 % enriched ^{87}Rb with 10 mbar of Ar. The atoms are heated to 97 °C. The color indicates the polarization. The level-scheme insets show the involved sublevels for each manifold of transitions. They are roughly ordered by frequency and OD. Allowed transitions are represented by solid arrows while singly and doubly forbidden ones by dashed and dotted arrows, respectively.

In order to improve the performance of the electromagnet, I applied two modifications compared to how AH previously operated it. Firstly, I switched the coils to be operated in series rather than in parallel, thus ensuring that exactly the same current flows through both coils, even if they have slightly different temperatures due to imperfect cooling. Secondly, I reduced the spacing between the ferromagnetic cores to the minimal possible distance, dictated by the physical dimensions of the vapor cell I was going to use, being 10 mm. With the nose-cones and a current of 20 A, the maximum the power supply can provide, I have measured magnetic fields of up to 2.36(4) T with the digital teslameter (Group3 DTM-151 with MPT-141 probe). For our purposes, however, the goal is not to achieve the highest possible field, but the most homogeneous one. By removing the nose-cones, the ferromagnetic cores have a cylindrical shape with a cross-section of 10 cm. By further reducing the spacing of the ferromagnetic cores to 10 mm again, the ratio between the distance and the diameter of the cylinders gets as close as possible to the approximation of an infinite plane. This configuration should give us the most homogeneous field over the longitudinal dimension of the cell as possible. In this configuration I measured a field of up to 1.75(4) T. This corresponds to a ground-state splitting of 54.3 GHz (for $m_I = +\frac{3}{2}$) and translates roughly to a frequency difference of 21 GHz between the outermost allowed π transitions of the spectrum.

Before performing the measurements presented in this and the next chapter I characterized the transversal homogeneity of the static magnetic field. The used procedure and our conclusions can be found in Appendix B. For a magnetic field of 784.43 mT a transverse homogeneity, averaged over the full longitudinal dimension of the vapor cell, of 27 µT over a distance of 5 mm is reached. Considering that the light beams are going to address much smaller regions of the cell, the interrogated atoms should experience the same magnetic field with spatial variations < 5 µT. According to AH's thesis [168] the magnetic field generated by this specific electromagnet drifts on the order of hundred µT on a 10 s to minutes timescale. For the transitions of interest for the memory scheme this translates to < 2 MHz shifts for a magnetic field of about 1.06 T. However, I could not observe such fluctuations. In fact, even over the course of several hours no significant drifts of the transition frequencies could be measured with the wavelength meter. Depending on the exact transition that was monitored, the shift expected from these field strength fluctuations is at the limit of what the wavelength meter can accurately resolve. Such fluctuations, however, would be in conflict with the spatial homogeneity measurement I performed. The imaging of the static magnetic field relies on capturing a sequence of images. The whole sequence takes several minutes, meaning that the measured homogeneity is averaged over the whole acquisition time. Given the measured spatial homogeneity, the presence of such temporal fluctuations does not seem plausible. Probably, re-wiring the magnet in series has helped in significantly reducing the temporal fluctuations.

Throughout the course of the experiments I repeatedly encountered issues with the water cooling of the electromagnet and its power supply. According to the manufacturer, the cooling circuits of these two devices must be connected in parallel.

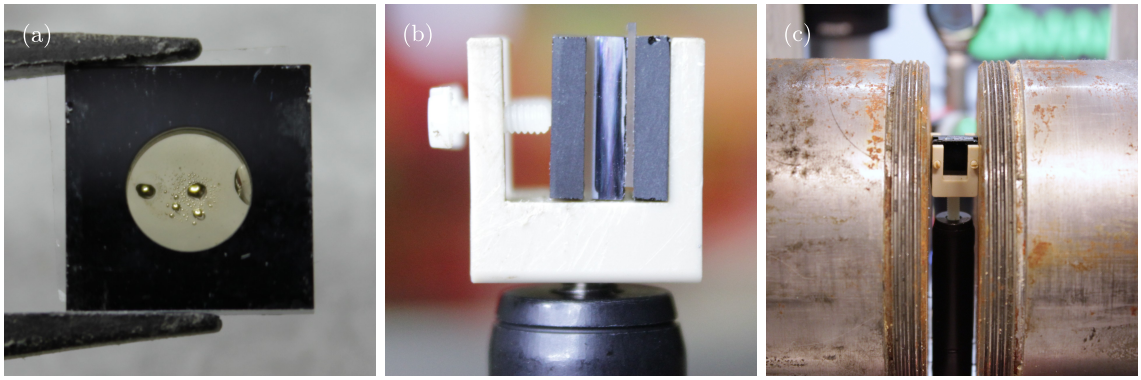

Figure 4.6: (a) Front view of the micro-fabricated vapor cell. Some Rb droplets are visible on the cell window. For cleaning the window, to improve transmission, I followed the 'cleaning instruction' from Appendix F of [168]. **(b)** Side view of the vapor cell sandwiched between two Schott RG9 filter glass plates in a PEEK mount on an aluminum post. **(c)** Vapor cell between the ferromagnetic cores of the electromagnet. The limited optical access due to the tight geometry imposed by the ferromagnetic cores becomes evident.

The clogging up over time of the power supply's heat exchanger with rust and limescale, however, led to changing cooling water ratios between the two devices. This caused changes in the thermal load of the coils, which in turn caused changes in the intensity of the magnetic field. This effect acts on a much longer timescales – about once a year the power supply started overheating, not receiving enough water to be properly cooled. By reversing the flow of the cooling water several times it was possible to restore the device's operability. In May 2022 the water cooling stopped working completely. Luckily, Dominik Sifrig, a technician in our department was able to repair the device by disassembling it, mechanically unclogging the heat exchangers and replacing old internal tubing. After this maneuver, we observed an abrupt change in the measured transition frequencies, which are compatible with a change of about 24 mT in the magnetic field strength.

4.4.2 Vapor Cell

Having in mind the homogeneity of the magnetic field along the optical axis we opted against a conventional (glass-blown) spectroscopy cell for this setup. We want to have a vapor cell that is thin compared to the dimensions of the magnet, so that all the atoms along the optical axis experience the same field strength. Furthermore, having a cell with small transverse spatial dimensions as well, allows us to move the ferromagnetic cores of the electromagnet close together, reaching the configuration described earlier, where the spacing between the cores is much smaller than their diameter.

A cell filled with isotopically enriched ^{87}Rb is advantageous for our purposes. Not only does it contribute to a higher OD for a given temperature at a fixed geometry, but it also removes the unwanted ^{85}Rb component to which the light could couple

off-resonantly, leading to noise in the memory setup. Additionally, having a buffer gas in the vapor cell brings the benefit of localizing the Rb atoms and limiting their motion from a ballistic to a diffusive one, which will be essential for the operation as a quantum memory. This strongly reduces the Rb-Rb collisions and leads to less wall collisions and the consequently induced depolarization of the atoms [233]. Choosing a low buffer gas pressure (a few tens of mbar), compared to what is commonly used for atomic clocks in such miniaturized cells, brings these advantages without excessively broadening the optical transitions.

For the experiments in the HPB regime performed for this thesis I thus used a micro-fabricated vapor cell, labeled as #106, made by Gaetano Mileti's group in Neuchâtel. The cell design was engineered in order to miniaturize atomic clocks [234] and was also successfully used for microwave imaging experiments [235]. For the mentioned applications, high buffer gas pressures are usually preferred/required. Fortunately, in Neuchâtel they also had an old cell with low buffer gas pressure we could use. The cell is composed of a silicon layer, in which a through hole is etched, and of two borosilicate plates. For achieving a hermetic seal, the glass is anodically bonded to the silicon substrate. These cells can be produced in batches on silicon wafers, but are filled and sealed one by one. A detailed description of the fabrication process of the cell type used here can be found in [236, 234] while for a more general review on MEMS cell fabrication I point the reader to [52].

Figure 4.6 shows the front and side view of the vapor cell as well as the confined space once the cell is placed in the electromagnet. The external size of the cell is $10 \times 10\,\mathrm{mm}^2$ with a thickness of 3 mm. The etched hole has a diameter of 5 mm and the internal thickness is 2 mm. The cell is filled with enriched ^{87}Rb of a purity of $\sim 90\%$ and about 11 mbar of Ar buffer gas (for more details about the reverse engineering process of determining the cell filling see Appendix C).

Especially for thin vapor cells, anti-relaxation wall coatings are of interest. The high temperatures of about 300 °C and the required voltages of up to kV for anodic bonding generally destroy these coatings. However, first successes were achieved in realizing micro-fabricated cells with anti-relaxation wall coatings by using a low-temperature bonding technique [237]. The utilized thin-film indium bonding technique only requires temperatures of 140 °C, which is low enough not to deteriorate wall coatings based on organosilanes, such as octadecyltrichlorosilane (OTS). Besides alkenes and alkanes, organosilanes are known for having good anti-relaxation properties for alkali metals. This coating is shown to preserve the alkali spin-polarization for up to 2000 wall collisions, while being more temperature resistant than alkenes and alkanes, which are too delicate even for the low temperature bonding techniques. A severe drawback of the kind of cells presented in [237] is that they lose atoms over time, drastically reducing the 'lifespan' of the cell (in the order of weeks) compared to ones sealed with anodic bonding. However, many low temperature bonding techniques are being researched, as room temperature anodic bonding [238], which may be compatible with wall coatings and be well suited to produce long-lived hermetic sealed cells.

4.4.3 Temperature Control by Laser Heating

According to [130] an OD > 5 is necessary to achieve a maximum total (or internal) efficiency above 30 % with forward retrieval for the physical process. Using a 2 mm thin vapor cell requires that it be heated up to 100 °C in order to reach an OD of 2.6 on the probe transition for an unpumped vapor. In order to avoid placing any magnetic materials within the high magnetic field, and due to the spatial confinement, we refrain from using resistive heating elements. Instead, for heating the cell we rely on high power infrared lasers and absorptive long-pass filters. This technique has already been successfully used in micro-fabricated cell devices such as magnetometers [239] and for frequency-tunable microwave field detectors [240].

We use two low-cost, fiber-pigtailed, multimode lasers (Seminex 4PN-108) emitting up to 3.3 W each at $\lambda = 1550$ nm[3]. The operating current is provided through laboratory power supplies (EA-PS 2016-100). For monitoring the laser diodes' temperature I glued PT100 thermistors to the diode package with thermally conductive epoxy; the temperature is read out with a digital temperature monitor (Jumo di eco). To ensure the cooling of the lasers they are mounted on Peltier elements whose hot side is thermally contacted to a large heatsink cooled by a fan. Additionally, a water cooler (originally intended for CPUs) is clamped on top of the diode packages to further cool the devices. The two laser beams are outcoupled on opposite sides of the vapor cell with diameters slightly larger than the cell's aperture. Due to the limited optical access, the beams could not be aligned on the experimental optical axis. Instead they are incident under an angle, coming up from the bottom (the setup is enclosed in a black foamboard box to block stray light). The path of the heating beam sets a limit on how close optical components can be placed to the cell of about 250 mm on either side[4]. The vapor cell itself is sandwiched between two 2 mm thick pieces of RG9 Schott filter glass (see Fig. 4.6(b)). This filter has an internal transmission of 94.67 % at 780 nm while being highly absorptive at 1550 nm (nominal transmission only 7.81 %)[5]. With an emitted power of ~ 800 mW per heating laser, the vapor cell reaches a temperature of about 93(1) °C. The temperature control approach presented here can be considered to be 'brute force', as it could be rendered much more efficient. In [239], for example, an interior volume of $(1.5\,\text{mm})^3$ is heated to 150 °C with only 140 mW using the same technique. To achieve this efficiency the vapor cell is thermally insulated from the environment by being placed in a second vacuum.

By directly heating the cell windows with this technique, the build-up of Rb condensation on the glass is highly suppressed since they are the hottest spot of the

[3] For my first measurements I actually used a Seminex 15P-112 laser, the same as in [168]. With an output of > 4 W the beam was split and distributed below the main optical axis in order to illuminate the cell from both sides. Finally, when it broke, I chose to replace it with two lasers in order to reduce the IR optics around the electromagnet.

[4] Optical components can be precariously mounted from above, but this is not a great idea if mechanical stability is a concern.

[5] Data sheet for RG9 Schott filter glass available at https://www.schott.com/shop/advanced-optics/en/Matt-Filter-Plates/RG9/c/glass-RG9.

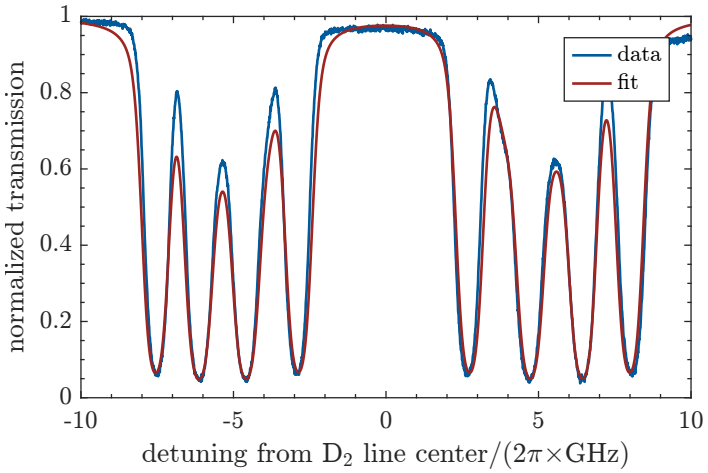

Figure 4.7: Exemplary fit with the software ElecSus to determine the temperature of the vapor cell. The measured π-polarized spectrum (blue) for a heating laser current of 2.7 A is shown. The frequency detuning is calibrated with the help of an absolute atomic frequency reference. In red the calculated spectrum resulting from the fit with some manual adjustments to best represent the width and depth of the transmission is shown. The resulting parameters are: $T = 97\,^{\circ}\mathrm{C}$, 35 MHz shift and 150 MHz additional broadening induced by the buffer gas, 10 % ^{85}Rb abundance, $B_{\mathrm{DC}} = 1.1\,\mathrm{T}$ with $B_\theta = 270°$.

cell. The Rb tends to deposit on the interior silicon layer. The thermalization time of this heating method for the previously described vapor cell is on the order of ten minutes.

A minor disadvantage of the RG9 filter glass is its opaque appearance at visible wavelengths, making the alignment of the laser beams to the cell center harder. Furthermore, the losses due to the filter glass, given by the high but not unity 'internal' transmission at 780 nm and the partial reflection on the filter surfaces, can become critical when working with single photons as in the scenario of the quantum memory. While for the proof-of-principle demonstration presented here they might be acceptable, in the future, optimization possibilities should be investigated, as e.g. anti-reflection (AR) coating the filters and using index matching gel between the filter and the cell windows. Ideally, one would fabricate a vapor cell where the windows are directly made out of the filter glass. An improved setup could also use a one-sided heating scheme with asymmetric filter thicknesses, where the front filter glass is thinner and absorbs about 50 % of the heating light and the second filter absorbs the residual light [239]. In this scenario the filter thicknesses should be designed such that the same amount of optical power is absorbed by each filter. This scheme would also reduce the mid-infrared optics to one side of the setup.

For determining the temperature of the vapor we directly interrogate the atoms. To record unperturbed spectra, we work in the weak probe regime. In order to still have a good signal to noise ratio for this type of measurement we use a collimated beam with a diameter of about 1 mm to probe the atoms. A Doppler-free saturation

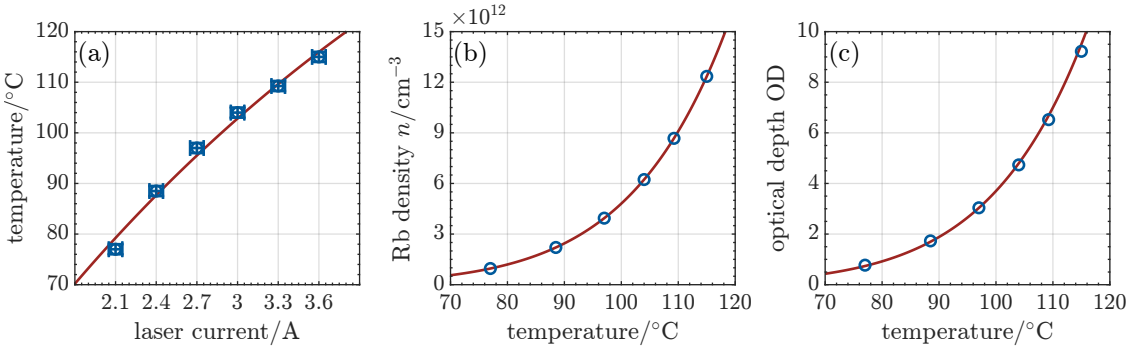

Figure 4.8: Calibration of the vapor cell temperature. **(a)** Atomic temperature extracted from the ElecSus fits as a function of heating laser current. The data points are fit with the function $f(x) = a(1 - \exp(bx))$. The uncertainty on the temperature is estimated to be $\pm 1\,^{\circ}\mathrm{C}$ while for the current $\pm 50\,\mathrm{mA}$ **(b)** The Rb density is modeled according to the equation in [228]. The points correspond to the densities for the measured temperatures. **(c)** The OD of the strongest π transition, computed from the transmission yielded by the ElecSus fit, is represented by the points. The curve is obtained from the model in (b) exploiting OD $\propto n$.

spectrum is recorded in parallel and used as an absolute frequency reference (details see Appendix D). For the measurements presented in this section a power in the order of $30\,\mathrm{nW} < I_{\mathrm{sat}}/100$ was used. A fraction of the laser beam is picked up before the vapor cell to be used as a power reference and to correct for frequency dependent power variations during the scan. The spectra are then fitted with the program ElecSus [229, 231]. The known cell parameters, as the cell's length, the shift and the additional broadening induced by the buffer gas are kept fixed. The first times we used the ^{87}Rb abundance as a fit parameter, but started to use the obtained value of about 90 % as fixed parameter since it consistently resulted in good fit results. Finally, the temperature value was fine-tuned manually to achieve the best match with the width and the flatness of the absorption curve. In Fig. 4.7 the measured spectrum for a current of 2.7 A per heating laser can be seen compared to the ElecSus fit. In ElecSus, the frequency detuning is defined relatively to the center of the line, weighted by the natural abundance of each isotope. Zero detuning thus corresponds to $\nu_{cl}^{\mathrm{D}_1} = 377\,107\,407.299\,\mathrm{MHz}$ and $\nu_{cl}^{\mathrm{D}_2} = 384\,230\,428.1\,\mathrm{MHz}$ respectively [229]. In order to avoid confusion I adopted this definition for the center of line frequency in this chapter.

To characterize the laser heating we recorded the transmission spectrum (of the π-polarized transitions) of the vapor cell at different heating laser currents. The resulting current dependence of the temperature is plotted in Fig. 4.8(a). The Rb density (for both isotopes) is modeled via the formula in [228]. Its dependency with respect to temperature can be seen in Fig. 4.8(b). The unpumped OD for the strongest π transition ($|+\frac{1}{2}, +\frac{1}{2}\rangle \longleftrightarrow |+\frac{1}{2}, +\frac{1}{2}\rangle$) is calculated directly from transmission obtained from the ElecSus model and shown in Fig. 4.8(c). The curve is obtained from the density curve in (b) by exploiting the proportionality relation between OD and number density. The points in (b) and (c) are not directly measured

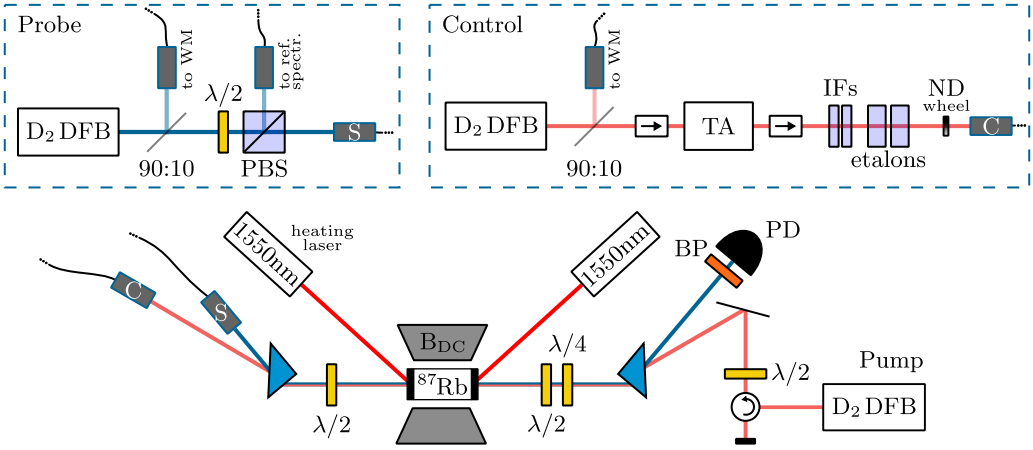

Figure 4.9: Experimental setup for the absorption spectroscopy measurements. The light fields are prepared in modular stages and are brought to the experiment through optical fibers. Probe and control beams are combined on a calcite prism before the vapor cell and are separated by the same means afterwards. The spectra are recorded on a PD. The pump is aligned in the counter-propagating direction and is coupled in through a circulator. The labels S and C represent the fiber connections for the probe (signal later on) and the control respectively. Used abbreviations: 90:10 – beam sampler with the specified branching ratio; WM – wavelength meter; PBS – polarizing beam splitter; TA – tapered amplifier; IF – interference filter; ND – neutral density gradient filter wheel; BP – 780 nm band-pass filter; PD – photodiode.

experimentally but derived from the temperature values obtained through the fit.

4.4.4 Spectroscopy Setup

The experimental setup used for the measurements presented in this chapter is sketched in Fig. 4.9. Three main lasers are used for the characterizations presented here: a probe, a pump, and a control amplified by a tapered amplifier (TA, TApro 780 nm with TA-0780-3000-2 chip). A set of IFs and etalons is set up downstream of the TA for filtering amplified spontaneous emission (ASE). The spectral cleaning of the control – indispensable for single-photon operation – will be discussed in detail in the next chapter (see Section 5.2.4). Towards the end of my work all lasers used in the setup were DFB diodes. Two models (Eagleyard EYP-DFB-0780-00080-1500-BFW01-0002 and -0005) were employed. According to the manufacturer's specifications they differ in terms of mode-hop free tuning range and linewidth, being 1500 pm and 2 MHz, and 25 pm and 0.6 MHz, respectively. For the Paschen-Back measurements presented in this chapter and the memory experiments in the following one, all the NIR lasers are operated free-running. The huge energy splitting in a large magnetic field makes it more complicated to lock the frequency of the lasers (see end of Ch. 5). The DFB lasers can be tuned in frequency by changing the device's temperature or its operating current. These lasers have a current tuning

coefficient of $0.003\,\mathrm{nm\,mA^{-1}}$. When paired with the digital laser controller (Koheron CTL200-1-B-200), which has a current resolution of $5\,\mu\mathrm{A}$, the frequency can be tuned in steps of about 7 MHz. The temperature tuning coefficient, on the other hand, is $0.06\,\mathrm{nm\,K^{-1}}$. Steps of one mK would correspond to frequency differences larger by almost an order of magnitude. The temperature tuning resolution ultimately depends also on the temperature stabilization of the DFB diode. For these reasons I usually rely on a current scan when scanning the laser frequency, considering as well the faster response time of the laser to current changes, and adjust the temperature to set the coarse center frequency of the scan. The frequency of the various lasers can be referenced with a wavelength meter (HighFinesse WS-7) through dedicated auxiliary fiber ports. A fraction of the probe light is branched off after the laser and sent to a Doppler-free saturation spectroscopy setup for absolute frequency reference (see Appendix D). The probe and control beams are combined on an AR coated polarizing calcite prism, while the pump beam is coupled in through a circulator in the counter-propagating direction following the control's path. This configuration results in the pump and control fields having the same polarization, orthogonal to that of the probe. The beams are tightly focused on the cell to achieve high control Rabi frequencies. A neutral density (ND) gradient wheel is used to vary the optical power of the control beam. The vapor cell, mounted in a custom-made PEEK holder, is placed in the static magnetic field generated by an electromagnet. It is heated with a mid-infrared laser and special absorbing filters surrounding the cell. For detection, a PD with a bandpass filter to suppress ambient light is used after probe and control beam have been separated by a second polarizing prism.

4.5 Results

The HPB regime has been accurately explored with precise spectroscopic measurements that were successfully compared to theoretical models with minimal deviations [220, 241, 242]. Several experiments in nm-thick vapor cells have been performed as well to study the behavior of alkali atoms in high external magnetic fields [223, 232]. In [243] the process of optical pumping in the HPB regime has been modeled and measured. For the measurements I present in the following, my intent was to familiarize with the high magnetic field regime before starting to implement the quantum memory, rather than doing precision measurements.

4.5.1 Absorption Spectroscopy

I started by characterizing the absorption spectrum of the vapor cell in the HPB regime. Figure 4.10(a) shows the D_2 line spectrum of the micro-fabricated vapor cell at $92(1)\,^{\circ}\mathrm{C}$ in an external magnetic field of 1.06 T. In order to achieve a good SNR, a collimated beam with a diameter of approximately 1.6 mm was used. The eight central peaks (blue) correspond to the allowed π-polarized transitions, probed with a horizontally polarized laser. Of these transitions, the outmost two belong to the

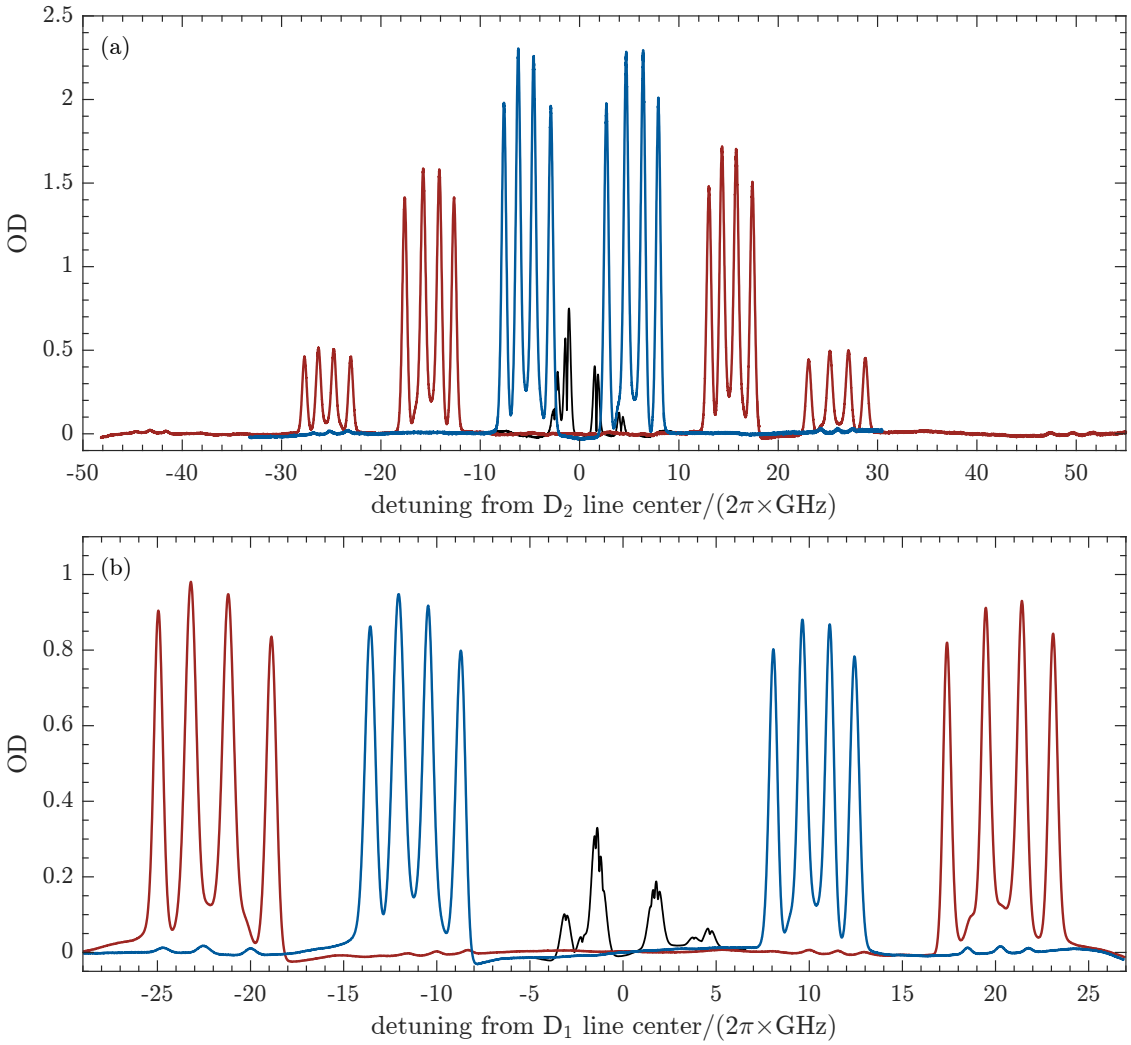

Figure 4.10: Spectroscopy of the MEMS cell #106 in an external magnetic field of 1.06 T at a temperature of 92(1) °C. Horizontally (blue) and vertically (red) linearly polarized light in the laboratory frame corresponds to π polarization and an equal superposition of σ^+ and σ^- polarization respectively. The (allowed) σ^+ (σ^-) transitions have a frequency detuning larger (smaller) than zero with respect to the line center. The strongest forbidden transitions can be identified in the spectra. The spectrum of a reference cell filled with natural Rb is plotted as black line to give a sense of scale, its amplitude is scaled arbitrarily. **(a)** D_2 line. **(b)** D_1 line. Note the different ranges of the x-axes.

$m_I = +\frac{3}{2}$ manifold, while for each step closer to the center of the spectrum the value of m_I decreases by one. The eight transitions (red) with the lowest frequency correspond to the allowed σ^- transitions, while those eight with the highest frequency (red) are the σ^+ transitions. These 16 transitions are all driven by the same vertically polarized probe beam. Furthermore, the strongest singly forbidden transitions can be identified in the spectrum: both π-polarized, the low frequency σ^- and the high frequency σ^+ triplets are visible. The displayed spectrum has been stitched together from three separate measurements for each laser polarization. Since the depicted frequency range exceeds the scanning range of the DFB laser that is achievable by current modulation alone, several measurements at different operating temperatures have been recorded to cover the whole range. The stitching has been done manually, since this is the only measurement for which scanning the current alone was not enough. If necessary, the laser diode controller used for these measurements (Koheron CTL200-1-B-200) could be programmed over the serial interface to do an automated stitched scan by changing the current and temperature, instead of using the simple graphic user interface provided by the manufacturer. Linear horizontal and vertical polarization have been recorded separately in two steps. The probe field is highly horizontally polarized by the calcite prism. For recording the vertical polarization, a half-wave plate is used to rotate the probe's polarization before the vapor cell. The spectrum of a reference cell filled with natural Rb, used as an absolute frequency reference (see Appendix D), is shown as black line to give a sense of scale to the recorded spectrum. To correct for frequency-dependent power variations of the laser I recorded the transmission of the probe through exactly the same setup without heating the vapor cell. For the short cell the effect of the atoms is negligible to a good approximation at room temperature. For the horizontal polarization this correction worked well, while for the vertical one I additionally modeled the remaining background with a third-order polynomial for correction. The measured spectrum is in good qualitative agreement with the computed spectrum and the work of [220].

The corresponding measurement was performed under the same conditions for the D_1 line as well. Figure 4.10(b) shows the spectrum recorded with a commercial 795 nm DFB laser (TEM Lasy 795). The D_1 laser is coupled to the experiment through the fiber port previously used for the D_2 probe. The half-wave plate is substituted with one for the correct wavelength. The spectrum shows the eight allowed π-polarized transitions (blue) and the two sets of four transitions for the σ^+ and σ^- transitions (red). All the singly forbidden transitions arising from the residual coupling of the ground state can be seen: two π triplets and one triplet for each σ polarization. Since the D_1 spectrum covers a smaller frequency range it could be recorded with one scan per polarization. Here the qualitative agreement with the expectations is good.

From here on the measurements in this chapter will focus on the D_2 line, since the next chapter will describe the implementation of a quantum memory at 780 nm. The (experimental) advantages of a pump laser on the D_1 line revealed themselves only at the very end of my time in the laboratory (for more details see Section 5.4.1).

4.5.2 From EIT to Autler-Townes Splitting

EIT has been extensively studied in the absence of external magnetic fields [151]. However, the HPB regime offers the possibility to isolate specific energy levels by inducing splittings larger than the Doppler broadened lines. To my best knowledge, only EIT in a ladder- [244] and in a V-scheme between three different fine-structure levels [245] have been studied in the HPB regime in warm atomic ensembles at the time of writing. In [246] EIT in a lambda-scheme in warm ^{85}Rb was investigated. However, with an applied magnetic field of 170 mT I would consider it to be an intermediate regime rather than the HPB regime.

Recalling Ch. 2, EIT originates from an interference between different atomic transition pathways. The coherently induced transparency can achieve complete transmission even in the weak control limit. Without this interference the transparency is simply given by a doublet structure of the absorption profile. In fact, the strong-field driven interaction leads to the splitting of the energy levels through the dynamical Stark effect, or dressing. In the so called Autler-Townes (AT) splitting regime the dressed-state frequency separation is proportional to the control Rabi frequency Ω [247]. For large Ω the interference becomes negligible [248]. The absorption profile transitions smoothly from one regime to the other. While in the strong-coupling regime (AT splitting) the transparency arises due to the gap between two resonances (two separate absorption lines), in the weak-coupling scenario (EIT) the reduction of absorption originates from the destructive interference of two absorption lines (difference between a broad and a narrow line) [249, 250].

I studied EIT in a lambda-scheme in the HPB regime because it is of interest for quantum memories, since the coherence between the states $|g\rangle$ and $|s\rangle$ can be long-lived when they are chosen to be in the $5\,^2\mathrm{S}_{1/2}$ ground state, avoiding to be limited by any excited state lifetime. Since our goal is to build a broadband memory, in order to be efficient at storing, the transparency window should be of the same order as the bandwidth of the photons to be stored. Since the transparency window's bandwidth depends on the Rabi frequency Ω, I studied its dependence on the control field's power.

We considered a lambda-scheme within the $m_I = +\frac{3}{2}$ manifold where the probe and control fields have orthogonal polarizations in the lab frame (see Figure 4.11(b)). The amplified control laser is tuned to the σ^+ transition which couples to the $|m_J = -\frac{1}{2}, m_I = +\frac{3}{2}\rangle$ state of the $5\,^2\mathrm{S}_{1/2}$ term, while the probe closes the lambda-scheme through the π-polarized transition coupling to the stretched state. For this measurement the probe was scanned over all π transitions and had a power of tens of µW. The optical power of the control field is successively increased by turning an ND gradient wheel filter in order to analyze the transition from EIT to AT splitting and observe the increasing bandwidth of the transparency window. In order to achieve the spatial overlap crucial for EIT, I coupled both the probe and control beams into the same auxiliary PM fiber positioned after the vapor cell as an alignment aid. For this measurement the atoms were heated to approximately 90 °C and a static magnetic field of about 1.2 T was applied.

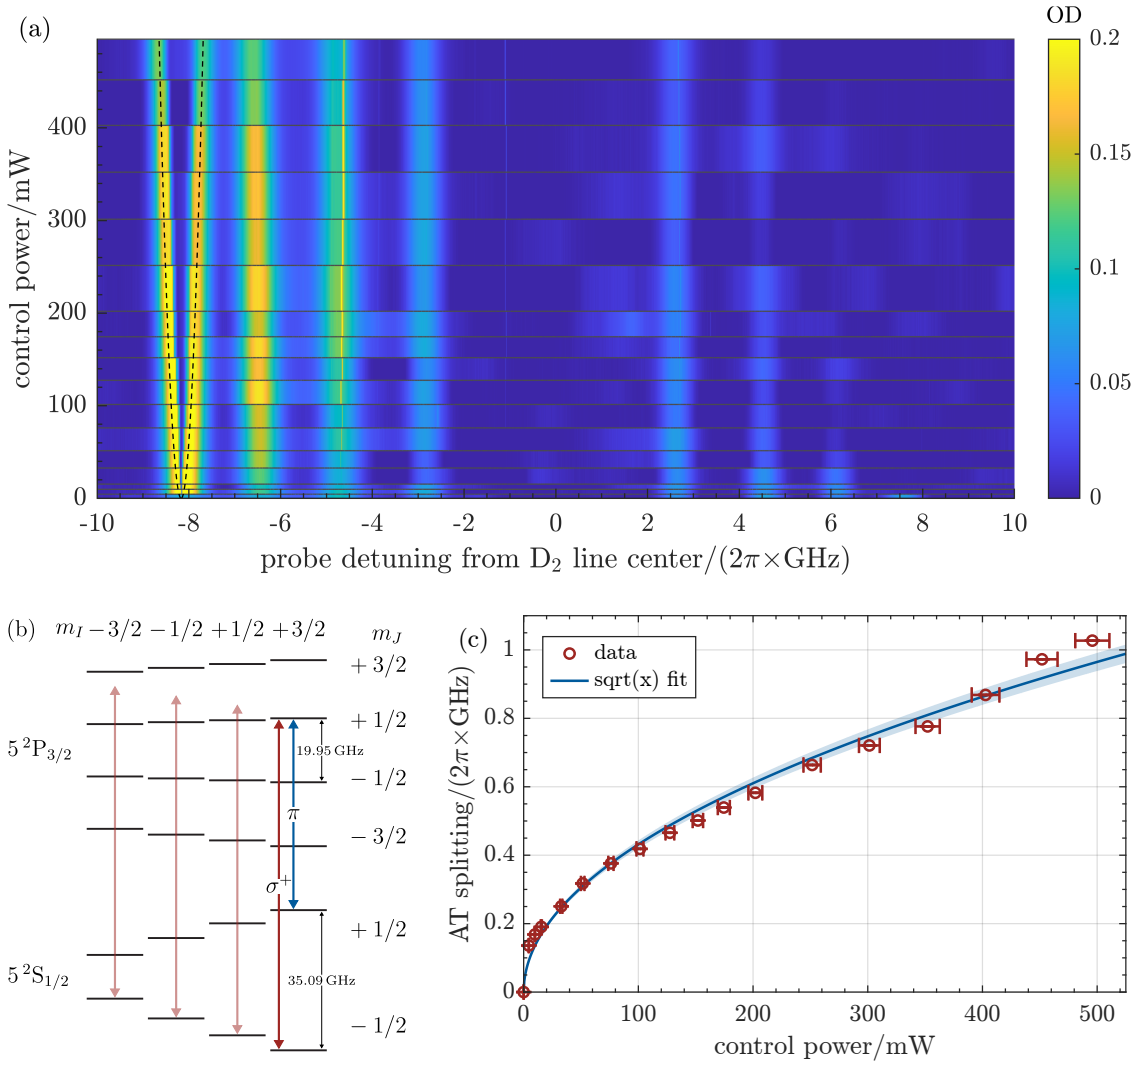

Figure 4.11: (a) Measured absorption of the probe beam as a function of the probe detuning for different control powers. The transition from the EIT to the AT regime can be observed. The maximum width of the splitting is approximately $2\pi \times 1\,\mathrm{GHz}$. The dashed curve shows the square root model obtained from panel (c) centered around the bare transition frequency. The horizontal lines represent the control power for which the data were measured. An alternative representation of these data is shown in Fig. E.1. **(b)** Level scheme for EIT in lambda-system. With higher Rabi frequency the off-resonant coupling of the control to the other $m_J = -\frac{1}{2}$ ground state can be observed in the spectrum (a) (see text). **(c)** Width of the transparency window as a function of the applied optical power. The data are fitted with a square root model. The shaded area represents the 95 % confidence interval of the plot.

Figure 4.11(a) shows the spectra for various control powers. For the lowest control power, the state it addresses already appears to be depleted. This is recognizable by the missing resonance on the far right of the spectrum. Its population is pumped into the stretched state. A narrow transparency window appears when the probe frequency is resonant to the second leg of the lambda scheme. When the control power is increased, this transparency window starts getting broader. The transition from EIT to AT splitting can be observed: the narrow transparency window evolves into two distinctly separate peaks. At the maximal power of 496(15) mW the transparency window reaches a width of about $2\pi \times 1\,\mathrm{GHz}$, which corresponds to the control Rabi frequency. It can also be noted that the control field starts to off-resonantly pump the adjacent ground state levels in the $m_J = -\frac{1}{2}$ multiplet to a significant degree. For control powers $> 150\,\mathrm{mW}$ the second resonance from the right vanishes as well. Furthermore, additional sub-Doppler wide features start appearing in the spectrum with increasing power. Most prominent are three (approximately) equispaced peaks. These features originate from the control field off-resonantly dressing the excited states with $m_J = +\frac{1}{2}$, which is further investigated in the next section. In the AT regime the frequency splitting, defined as the difference between the two maxima, corresponds to the control Rabi frequency Ω_c [247]. The Rabi frequency is a linear function of the electric field and can therefore be written as $\Omega_c \propto \sqrt{P}$, where P is the power of the control field [251]. Figure 4.11(c) shows the splitting as a function of the optical power fitted with a square root function. The model obtained from the fit is plotted as well in panel (a) as a dashed line. In the EIT regime the frequency splitting scales as the square of the Rabi frequency. Probably, the lowest control power levels used, corresponding to $\Omega \approx 2\pi \times 135\,\mathrm{MHz}$, are already too high to see this behavior in the measured data.

Avoided Crossings In order to identify the origin of the higher order features observed in the previous measurement, I analyzed their frequency dependence. With the same experimental parameters as in the previous measurement, choosing almost the maximum control Rabi frequency (about $2\pi \times 950\,\mathrm{MHz}$ corresponding to 485 mW), the control detuning was scanned over several GHz around the resonance frequency. The control frequency was measured with the wavelength meter. The resulting spectra are plotted in Fig. 4.12. The analyzed sub-Doppler peaks are part of avoided crossings (or anti-crossings), a typical signature of dressed states when the fields are detuned [247]. For zero control detuning the scenario of Fig. 4.11 is reproduced. The dressed states are at their minimum splitting for this control frequency. By detuning the control from the transition frequency, the splitting increases. In the asymptotic limit the larger peak of the doublet is located at the bare transition frequency to $|e, N\rangle$, here corresponding to the π transition coupling to the stretched state. The smaller peak, on the other hand, is asymptotically described by a line with slope 1 tending towards the energy of the uncoupled state $|g, N+1\rangle$. The amplitude of this secondary peak decreases for increasing control detuning. In the chosen control detuning range the control field becomes resonant with three out of four σ^+ transitions coupling to the ground state $m_J = -\frac{1}{2}$ manifold, closing the lambda-scheme with the scanned probe. In the figure we can thus

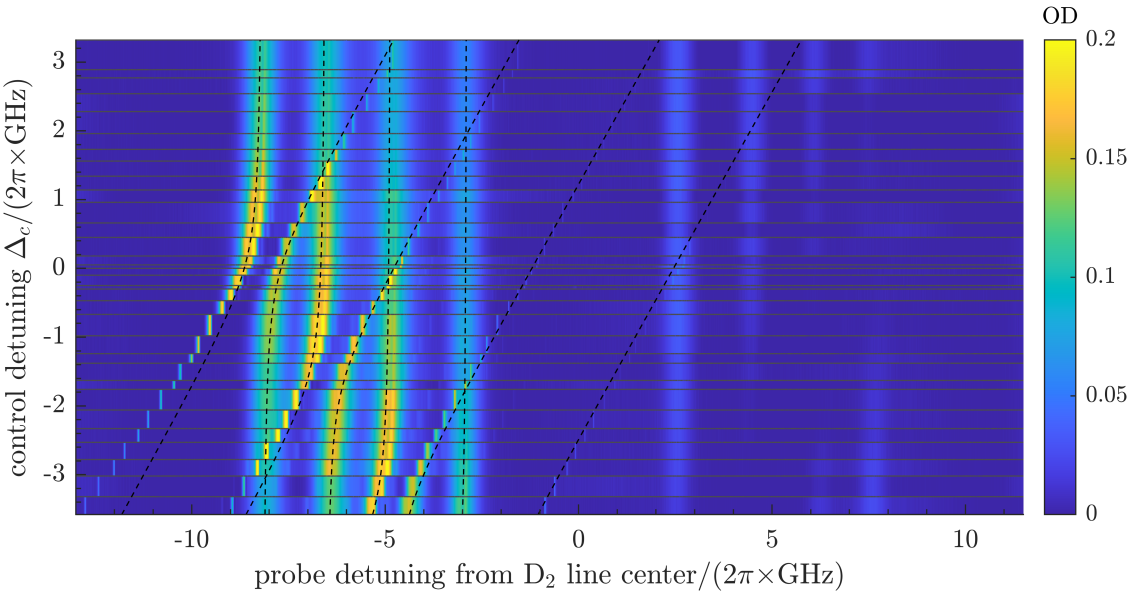

Figure 4.12: Measured absorption of the probe beam as a function of the probe detuning for different control detunings. Driving the atoms with almost the maximum control power results in dressed states. The structure of three avoided crossings can be observed, while the tail of a fourth one is recognizable. The detuning of the control frequency is chosen with respect to the σ^+ transition coupling to the $|m_J = -\frac{1}{2}, m_I = +\frac{3}{2}\rangle$ ground state. An alternative representation of these data is shown in Fig. E.2.

identify three avoided crossings and spot the asymptotic tail of the fourth one. For further red detuning of the control frequency the same behavior is expected for the transitions from the $m_J = +\frac{1}{2}$ ground states. At the given power, the control laser can off-resonantly dress all four $m_J = +\frac{1}{2}$ excited states simultaneously.

The dashed lines in the figure model the frequency of the coupled $|\phi_+\rangle$ and $|\phi_-\rangle$ states as a function of the detunings. Each avoided crossing is modeled as an ideal three-level system by [247]

$$\Delta_\pm = \frac{1}{2}\Delta_c \pm \frac{1}{2}\sqrt{\Delta_c^2 + \Omega_c^2} \tag{4.9}$$

describing the two branches of a hyperbola. The minimum splitting is achieved when the control field is tuned to resonance to the bare transition and is equal to Ω_c. In the model, I used a value $\Omega_c = 2\pi \times 950\,\text{MHz}$, and along the horizontal axis the hyperbolas are centered around the unperturbed transition frequencies of the probe. The y offset is adjusted to give the best overlap. While the horizontal axis is calibrated with a reference spectroscopy, it can be seen that the model curves and the data diverge on the low-frequency end of the axis. This behavior can be attributed to a non-linear scan of the probe laser, since all transitions used for the calibration are close to the center of the horizontal axis.

4.5.3 Polarizing the Nuclear Spin

For the quantum memory experiments described in Chapter 5 it is of paramount importance to prepare the initial atomic state in order to suppress spurious processes. Preferably, all atoms should be in the same state, with our target being the stretched state. Considering only allowed optical transitions, however, it is only possible to optically pump within one m_I manifold, since the optical transitions don't couple to the nuclear spin. If we assume a thermal distribution, with equal population in all eight ground states, only one quarter of the total number of atoms will be available. For a proof-of-principle demonstration this approach might be sufficient. Given the dependency of the storage efficiency on the OD, however, a more elaborated pumping scheme including all ground states should be investigated.

Another possibility that could be explored is using microwaves for preparing the initial atomic state. The strength of the microwave transitions inducing a change in the nuclear spin also decrease with a stronger applied field [168]. This sets the requirement for high microwave power levels. For the transitions between $m_J = -\frac{1}{2}$ and $m_J = +\frac{1}{2}$ ground states, this could pose a technical difficulty, as the required frequencies at the considered magnetic field strength are in the range of 25 – 35 GHz. For the transitions within one m_J manifold, where finding suited amplifiers should be no problem, an open question is whether these transitions can be driven individually or not, since the frequency differences are small. The use of microwaves for state preparation could be interesting, but has not been further investigated in the context of this work.

By optically pumping along an allowed transition, the population of the depleted ground level is transferred to the corresponding, other m_J-valued ground state within the same m_I manifold. Figure 4.15(a) shows such a pumped spectrum. For CW pumping, the $|m_J = -\frac{1}{2}, m_I = +\frac{3}{2}\rangle$ sublevel is fully depleted and m_I is conserved. In addition, the pump field couples off-resonantly to the adjacent level, weakly pumping that as well. I will go into more detail about the allowed state preparation in the next chapter. In Section 5.2.5 I will estimate the efficiency of the optical pumping process within one m_I manifold with a pump-probe experiment determining the atomic polarization.

In the calculated D_2 spectrum (see Fig. 4.4), we can identify the six triplets of singly forbidden transitions. Our main interest in these transitions is to drive them in order to help us pump the nuclear spin of the atoms. The two π polarized transition triplets have the strongest transition strengths. However, their frequencies overlap with those of other transitions. In order to avoid possible issues with driving some other transitions, even if they have the wrong polarization, I opted not to work with these transitions. The transition strength of the inner σ^+ and σ^- is so weak that they do not appear in the measured spectrum (see Fig. 4.10). This leaves the two outermost triplets in the spectrum for further investigation. For convenience I will number these transitions from 1 to 6 in the order of increasing frequency (as labeled in Fig. 4.4).

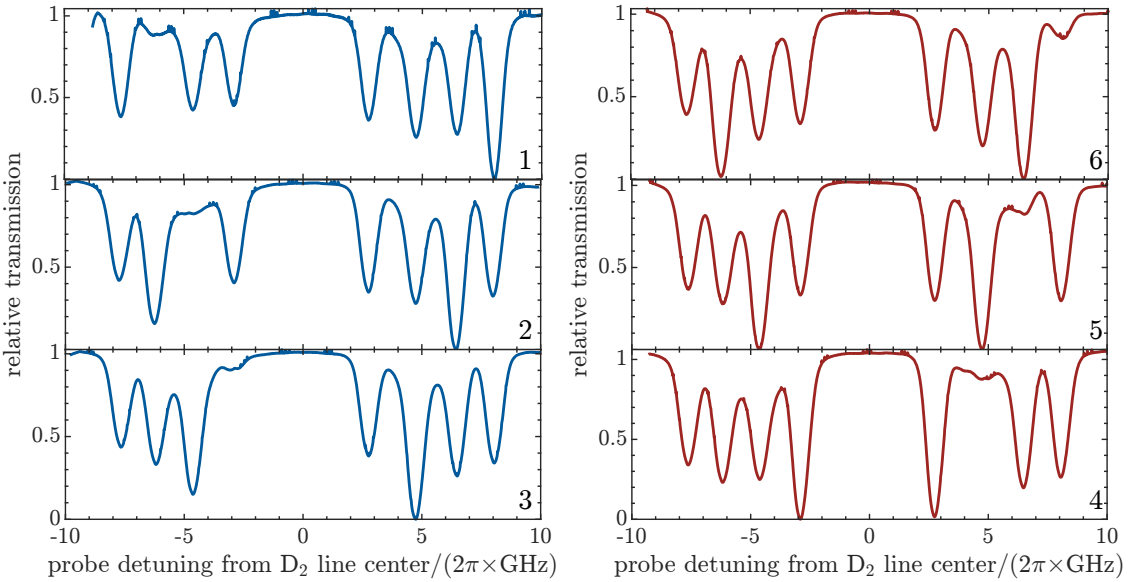

Figure 4.13: Pumping the nuclear spin by driving singly forbidden transitions. The measured spectra of the eight π-polarized transitions are recorded while the atoms are being continuously pumped. The pump frequency was tuned to be on resonance with the outmost forbidden transitions. The left-hand side plots (blue) show pumping on the σ^- transitions labeled 1 to 3, while the plots on the right-hand side (red) on the σ^+ transitions 4 to 6. The same nomenclature as in Fig. 4.4 is used to label the pump transitions. Some artifacts (three small recurring peaks) originating from the I/O card used for the EOM lock (see next chapter), set to full transmission here, are visible. The background is modeled as a third-order polynomial and corrected for.

In order to experimentally test whether it is feasible to polarize the nuclear spin by optical pumping with the available optical power, I performed a spectroscopy while the atoms where being continuously pumped. While the probe laser was scanning over all eight allowed π transitions, the pump laser was successively tuned to the forbidden σ^+ and σ^- transitions 1 to 6. The panels in Fig. 4.13 show the signal detected on the PD. The numbering of the panels corresponds to the forbidden transition that was driven by the pump light. We observe that the ground state level addressed by the pump is, as expected, highly depleted therefore rendering the vapor cell almost fully transmissive at the corresponding frequency. The pumped population appears to populate both states of the target m_I manifold. This manifests itself in the spectroscopy as increased absorption in peaks 1 and 8, 2 and 7, 3 and 6, or 4 and 5 respectively. The addressed excited state, however, has a total electronic angular momentum $m_J = \pm\frac{3}{2}$ (for $\sigma^\pm$-polarized forbidden pump) and thus cannot decay into both ground states. This suggests that the atoms are redistributed between the electron spin states by collisions. The forbidden pumping process is slower than the lifetime of the m_J states, but faster than the m_I lifetime. The nuclear polarization being longer lived than the rate at which the atoms are forbiddenly pumped opens up the possibility for optically pumping all the atoms into a single state.

This result motivated me to try to combine different allowed and singly forbidden transitions for pumping. With the memory experiment in mind, the desired initial atomic state would be the stretched ground state $|m_J = +\frac{3}{2}, m_I = +\frac{3}{2}\rangle$. By using the forbidden σ^- transitions with $\Delta m_I = +1$ (1 to 3) while simultaneously depleting the $m_J = -\frac{1}{2}$ states through allowed transitions, it should be possible, at least in principle, to fully transfer the atomic population into the desired target state.

In order to avoid unwanted coherences between the various required light fields, and in the worst case pump the atoms into some dark state, I decided not to use a phase modulator to generate sidebands for driving several transitions with only one laser. The recent developments in semiconductor lasers and their availability at the Rb wavelengths, as well as their affordability and the ready-to-use butterfly packaging, make us consider using one dedicated laser for driving each single transition. Especially, when paired with digital controllers for butterfly laser diodes (Koheron CTL200-1-B-200) the price and the footprint of the whole laser package is significantly reduced. By using fiber pigtailed butterfly diodes one could even eliminate the whole laser assembly from the optical table, freeing valuable space for the actual experiment.

Therefore, I decided to test this pumping scheme as far as possible with the lasers we had available in the laboratory. The pump laser was tuned to the allowed σ^+ transition coupling to the $|m_J = -\frac{1}{2}, m_I = +\frac{3}{2}\rangle$ ground state. I installed two further DFB butterfly lasers acting as forbidden pump and $m_I = +\frac{1}{2}$-pump respectively. They were aligned to the main pump beam under a small angle using D-shaped mirrors. The forbidden pump was tuned to the frequency of the singly forbidden σ^- transition coupling to the $|m_J = +\frac{1}{2}, m_I = +\frac{1}{2}\rangle$ sublevel (1 from before), with the scope of transferring the atomic population from the $m_I = +\frac{1}{2}$ to the $m_I = +\frac{3}{2}$

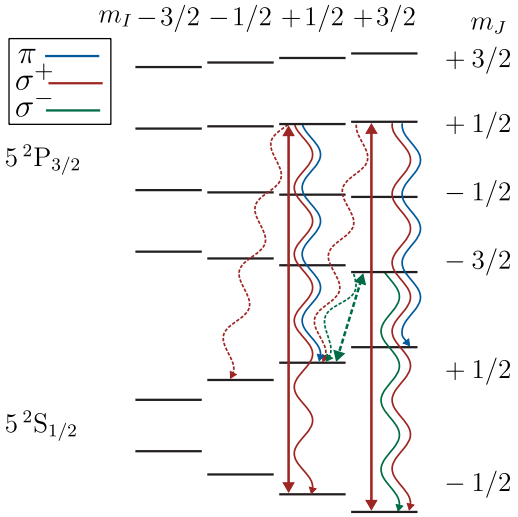

Figure 4.14: Energy-level diagram depicting the optical pumping scheme. Two σ^+ polarized lasers are used to pump the $m_I = +\frac{1}{2}$ and $m_I = +\frac{3}{2}$ ground state manifolds into the respective $m_J = +\frac{1}{2}$ state. A forbidden σ^- transition is used to (partially) polarize the nuclear spin. Allowed transitions are represented as solid lines, while forbidden ones as dashed lines. Undulated lines represent radiative decay channels. The color of the transitions identifies the polarization.

manifold. Simultaneously, the $m_I = +\frac{1}{2}$-pump, driving the allowed σ^+ polarized transition coupling to the $|m_J = -\frac{1}{2}, m_I = +\frac{1}{2}\rangle$ ground state, repopulates the state addressed by the forbidden pump. The pump scheme is visualized in the energy scheme in Fig. 4.14. All three pump lasers utilized for this measurement were operated at an optical power of 20 mW. Even though the DFB lasers could provide a maximum output power of about 100 mW, the selected power limit was chosen to be compatible with the optical switches used for the quantum memory implementation (see next chapter). For this measurement, the probe laser was again scanned over the allowed π-polarized transitions.

Figure 4.15 shows the spectra of different pumping scenarios (red) building up to the scheme just described. All panels include a reference trace of the unpumped vapor (blue) for comparison. Panel (a) shows the pumped spectra for the scenario in which only one of the pumping lasers is turned on. One can see that the corresponding ground state is depleted. In the panel (b) the interplay between any two of the three pump lasers is analyzed. Finally, Fig. 4.15(c) shows the interplay of all three pump lasers together. To correct for frequency dependencies of the emitted laser power, I fitted the background of the unpumped spectrum with a second-order polynomial which I then used to correct all shown traces. The green curve in panel (b) shows that the use of the forbidden pump already increases the absorption of the stretched state. The strongest absorption of the target state is achieved in panel (c). However, the absorption of the ground state $|m_J = +\frac{1}{2}, m_I = +\frac{1}{2}\rangle$, addressed by the forbidden pump, seems unchanged. The same behavior is already visible in the red trace in panel (b). This data corresponds to the scenario where the two pump lasers only address atoms initially with $m_I = +\frac{1}{2}$. This indicates the presence

of a non-radiative spin relaxation process, e.g., wall collisions, competing with the nuclear spin pumping and transferring atoms back to $m_I = +\frac{1}{2}$. From the recorded spectra the $|m_J = -\frac{1}{2}, m_I = +\frac{1}{2}\rangle$ ground state appears to be depleted, while for $|m_J = +\frac{1}{2}, m_I = +\frac{1}{2}\rangle$ no significant difference with respect to the reference curve can be identified. It can thus be estimated that 50% of the atomic population is moved from $m_I = +\frac{1}{2}$ to $m_I = +\frac{3}{2}$. The presented pumping scheme can be extended to all ground states. First, each m_I manifold is pumped into the respective $m_J = +\frac{1}{2}$ state. Subsequently, forbidden pump lasers are used to transfer the atomic population to higher m_I values. Assuming each manifold can be half depleted, and under the same experimental conditions, it should be possible to achieve a total atomic polarization of $\frac{15}{32}$, despite the isolation of nuclear states from one another.

A natural follow-up of this experiment would be to investigate how the atomic ensemble reacts to different power ratios between the pump lasers.

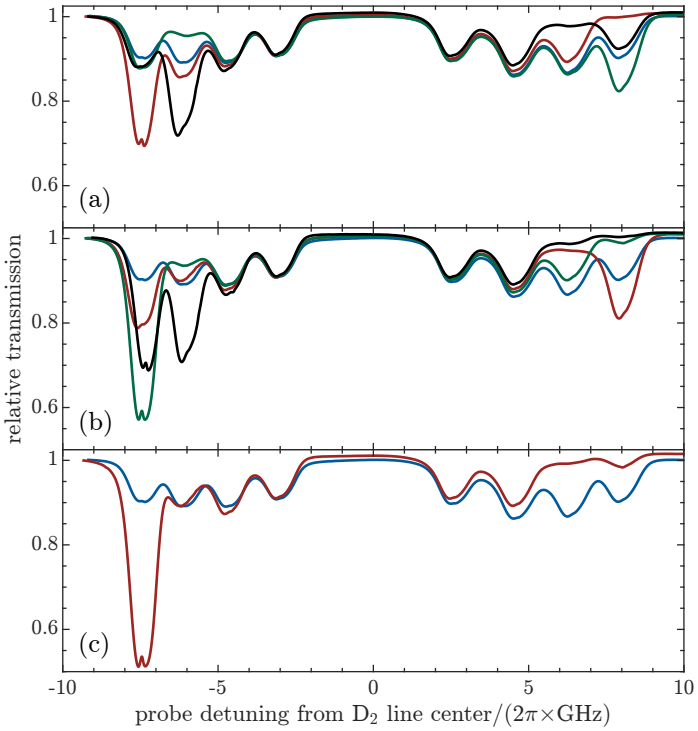

Figure 4.15: Combined optical pumping scheme involving allowed and singly forbidden transitions. The applied magnetic field was approximately 1.06 T. The unpumped spectra are shown in blue for comparison. The background of the unpumped spectrum is modeled by a second-order polynomial and is used to correct all plotted traces. **(a)** The effect of a single pump laser is shown, either: pumping along the allowed σ^+ transition (red), or: the forbidden pump drives the forbidden σ^- transition (green), or: the $m_I = +\frac{1}{2}$-pump drives the allowed σ^+ transition (black). **(b)** The combined effect of two lasers is shown. The forbidden pump and the $m_I = +\frac{1}{2}$-pump are shown in red, while the pump in combination with the forbidden pump is depicted in green. The effect of the pump and $m_I = +\frac{1}{2}$-pump is shown in black. **(c)** All three lasers, as shown in the scheme of Fig. 4.14, are used for pumping.

5 Quantum Memory in the HPB Regime

The first of the two approaches we pursue in our laboratory to isolate a clean three-level system for realizing a quantum memory is to exploit the HPB regime. The quantum memory scheme presented here works and is implemented on the D_2 line of ^{87}Rb. Corresponding approximately to the wavelength used to read and write CDs, this wavelength has the advantage that semiconductor lasers are easily available and cheap due to the industrial development. Furthermore, GaAs QDs emitting at 780 nm with excellent single photon characteristics have already been demonstrated (see Ch. 2). A further motivation for exploring the HPB regime is that the bandwidth of the photons an EIT memory can store is physically limited by the splitting between the two involved ground states. At zero applied magnetic field ^{87}Rb has a ground-state splitting of 6.8 GHz allowing for the storage of quantum dot photons with a bandwidth of about 1 GHz. Placing a quantum dot in a microcavity to enhance its extraction efficiency, broadens the emitted photon's bandwidth through the Purcell effect. For Purcell enhancement factors of over 10, end-to-end efficiencies of 57 % have been achieved [63]. However, the generated photons become too broad for a 'standard' Rb quantum memory. By operating a memory in the HPB regime the ground state splitting can be increased, shifting the bottleneck for the maximal bandwidth to the control pulse peak power and duration.

In this chapter I present the implementation of a broadband HPB quantum memory in a hot Rb vapor in a micro-fabricated vapor cell. As a proof-of-principle experiment and for exploring the vast interdependent parameter space in the attempt to find an optimal working point I first tested/operated the quantum memory with weak coherent pulses (WCPs). These results are published in [65]. At the beginning of this project it was not clear whether such a thin vapor cell would be at all suited for implementing a quantum memory. During the last weeks in the laboratory, before I started writing this thesis, motivated by the promising results achieved with the attenuated laser pulses, I performed first interfacing experiments of the HPB memory with the SPDC source presented in Ch. 3. For these experiments Gianni Buser helped me out by operating the heralded photon source and in general he assisted me in finding good operational parameters for the memory. First promising results were obtained, but the measured SNR were not yet high enough to show a non-classical retrieval signature. Improvements to the memory are ongoing while I am writing.

5.1 Memory Scheme in the HPB Regime

The HPB regime allows for the engineering of near-ideal multi-level systems in warm atomic ensembles. So far, non-degenerate three-level (such as ladder- [244] and V-scheme EIT [245]) and four-level systems (for a diamond FWM scheme [252]) between different electronic states were isolated to simplify the study and modeling of the atom-light interaction in these systems. In what is, to my best knowledge, a novel approach to quantum storage, in this thesis I present a lambda system, isolated exploiting the HPB regime, with both 'legs' resonant to transitions between the same two electronic states. Through the large energy splittings of the atomic sublevels, compared to the optical linewidth of the broadened hot vapor, and the decoupling of the nuclear spin **I** from the total angular momentum of the electron **J**, it is possible to optically address single atomic sublevels in a Doppler-broadened vapor. Such a degree of control is otherwise only possible in ultracold atomic systems at the expense of a much higher experimental complexity. The decoupling of **I** and **J** allows us to consider, to the first order, a single m_I manifold for the memory scheme, given that optical transitions don't couple to the nuclear spin.

When choosing the atomic sublevels for composing the lambda-scheme, the transitions strengths, the polarizations, and the energy splittings should be taken into account. Furthermore, it is advantageous to choose a stretched state as ground state $|g\rangle$. If all atoms are initially prepared with a fully aligned spin they cannot undergo spin-exchange collisions due to conservation of angular momentum. This technique works for arbitrary magnetic field strength, but cannot completely eliminate this type of relaxation mechanism [253].

For this reason we choose $|g\rangle = |m_J = +\frac{1}{2}, m_I = +\frac{3}{2}\rangle$ as the state to initially prepare the atoms into and consequently $|s\rangle = |m_J = -\frac{1}{2}, m_I = +\frac{3}{2}\rangle$ as the second ground state. These two ground states are both coupled through the excited state $|e\rangle = |m'_J = +\frac{1}{2}, m'_I = +\frac{3}{2}\rangle$ of the $5\,^2\mathrm{P}_{3/2}$ term, completing the lambda-scheme. The second stretched state, $|m_J = -\frac{1}{2}, m_I = -\frac{3}{2}\rangle$ would also be a valid choice for $|g\rangle$: the transitions strengths and the detunings to other transitions are similar to the system we chose, the ground-state splitting, however, is about 10 GHz smaller.

I chose to have the signal resonant to the π-polarized transition coupling to $|g\rangle$, in order to have the higher transition strength on the single photon signal leg of the lambda, whereas the (weaker) σ^+ transition is addressed by the amplified control beam. By working in the $m_I = +\frac{3}{2}$ manifold we take advantage of the largest energy splittings. The transitions involved in the memory scheme can be seen in the level scheme in Fig. 5.1. Operating the memory at a slight detuning Δ from the excited state $|e\rangle$, while maintaining signal and control in two-photon resonance, can be advantageous for discriminating noise at the natural atomic frequency.

The (partial) atomic state preparation is done with a dedicated pump laser on the same transition addressed by the control. A major consideration for the experiment design is to have a pump with the same polarization as the control so that it can be added to the control's optical path, avoiding to add any further optics

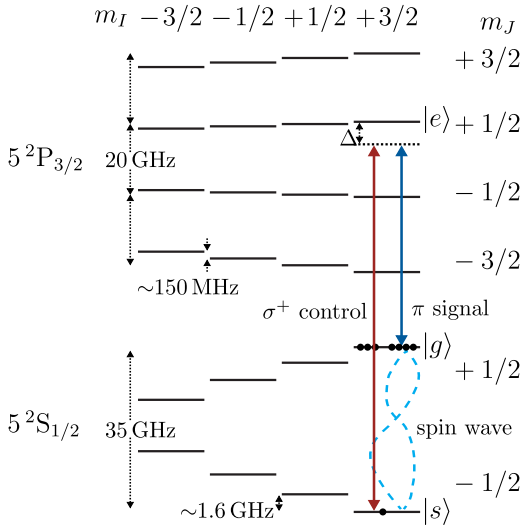

Figure 5.1: Memory scheme in the HPB regime. Ideally, all the atoms are prepared into one nuclear spin manifold. The stretched state is chosen to be $|g\rangle$. The signal transition is chosen to be π polarized for the higher transition strength. The memory is operated with a red-detuning Δ from $|e\rangle$ to be further detuned from the adjacent π polarized transition, while keeping two photon resonance. The shared coherent excitation mapped onto the atoms is split between the two m_J ground states of one m_I manifold.

to the signal arm. The σ^- polarized transition coupling to $|s\rangle$ is not suited for pumping, since, although being the stronger transition, the addressed excited level does not decay into $|g\rangle$. We also used the σ^- forbidden transition coupling to the $|m_J = +\frac{1}{2}, m_I = +\frac{1}{2}\rangle$ ground state (see Section 4.5.3) to further increase the OD on the signal transition.

5.2 Implementation with Weak Coherent Pulses

The parameter space of a hot vapor quantum memory is vast and interdependent. Before jumping directly into it with probabilistic single photons, it is recommendable to start with WCPs to find a good working point, and particularly in this scenario, to prove the feasibility of using such a short vapor cell as a quantum memory.

From the Zeeman memory we have learned how challenging it can be to switch from a deterministic mode of operation, as with WCPs, to a probabilistic one. For this reason I decided to directly set up the HPB memory so that it could react to an external trigger, even though the memory was first tested with attenuated laser pulses, which are deterministic. This way only little changes need to be performed to the setup for switching to a probabilistic source. The requirements a reactive memory implies for the insertion delay and switching times will be discussed later, when I describe the storage of SPDC photons.

I used exactly the same devices for the quantum memory described in this chapter as were used for [29], as they demonstrated to be well suited for building a reactive memory. The only exceptions are the arbitrary waveform generators (AWGs), where I used the hardware-upgraded version of the same model, and the single-photon detectors – an upgrade worth its own section later on. For a detailed account of the jitter contribution of electronic components, I redirect the reader to Ref. [254]. The main contribution to the jitter budget in that scenario was the SPAD.

The experiments described here are triggered by a digital delay generator (DDG), and not by a detector click. This leaves us with the jitter of the electronics and the EOMs which is small enough to not significantly affect the efficiency of the storage process. Jitter on the electronics triggering the control pulse generation has a negative effect on the efficiency of the storage process since it affects the time alignment of signal and control pulse.

All the lasers used for this experiment are CW DFB diodes. They are free running without any feedback on their frequency. Given the short duration of the experimental cycles, the frequency stability of a free running diode laser has proven to be sufficient. Using saturation spectroscopy in a reference cell to tune a laser to a specific transition, as is usually done in a no-field scenario, would be only possible if the reference cell itself was in the high magnetic field. Due to the geometric constraints, however, this was not an option in the described setup. For this reason, each laser has an auxiliary fiber link for measuring its frequency with a precise wavelength meter (HighFinesse WS-7). The frequencies are monitored throughout the experimental sequence, prioritizing the signal and control frequencies, and adjusted if necessary.

The experimental setup can be divided into different modules: the electronic logical pulses, the signal, the control, the atomic state preparation, the vapor cell comprising the core of the memory, the filter stage, and finally the detection stage. In the following sections I will focus on each single building block, explaining how they are implemented in this work and what their requirements are. The complete setup is sketched in Fig. 5.2.

5.2.1 Logical Trigger and Follower Pulses

For the generation of the logical signals used to time the experimental sequence we rely on the Highland Technology T564 digital delay- and pulse-train generator. This DDG impresses with its 20 ns short insertion delay and a RMS jitter of < 35 ps for an external trigger. It has a temporal resolution of 10 ps. The maximum repetition rate f_{rep} is determined by the duration of the programmed pattern t, with $f_{rep} = (t + 60\,\text{ns})^{-1}$ which only marginally affects the experiment rate. Furthermore, it can be operated in train mode where the DDG repeats the output of one or all channels after the programmed sequence has ended. The number of repetitions can be freely chosen by the user. The delay between the repetitions can be set to an integer multiple of 20 ns with a minimum delay of 60 ns.

The 4 channels of one device do not suffice for the generation of the logical signals required for the storage and retrieval sequence. Therefore we use two DDGs. When storing WCPs, the first DDG is triggered by its internal DDS synthesizer. This first device triggers the most time sensitive tasks: it switches off the pump beams by sending follower pulses to corresponding optical switches and triggers the generation of the WCP at the single photon level. Furthermore, it triggers the second DDG which relays the trigger to the control-pulse AWG, the time-tagger and a fast oscilloscope used to monitor the control pulse shape.

5.2.2 Optical Pulse Generation

For the operation of the broadband memory we need to be able to generate optical pulses with temporal widths of single ns. These pulses constitute our proxy signal and the control pulses responsible for writing and reading out the single photons to/from the memory. These pulses are first generated electrically with an AWG and then 'carved' into a CW laser field with an electro-optic amplitude modulator.

Arbitrary Waveform Generator The electronic pulse for driving the amplitude modulator is generated with a PicoQuant PPG512: an AWG designed to drive fast EOMs. It is composed of a programmable DAC and an amplifier and generates electrical pulse patterns that can be repeated in cyclic sequence. A 512 byte pattern with an 8 bit resolution can be programmed into the device's high-speed internal memory. This pattern is read out at 5 GS/s, resulting in 200 ps duration per byte and a total record length of 102.4 ns. The AWG has an insertion delay on the order of 10 ns. The specification of the jitter to one sample by the manufacturer is rather conservative. GB has measured it to be < 100 ps [254]. The output amplitude can be set between 0 and 12 V. For these experiments I used the revised hardware version of these devices, which has no limitations on the pulse length, other than the record length.

The pattern can either be played as a loop or it can be started and stopped by external trigger signals. The AWG has two modes of operation: 1) the 'reset' mode in which the device is free running and continuously repeats the programmed pattern for a low voltage trigger input (or no input at all). While for a high voltage trigger input, the output is interrupted at the current byte. 2) the 'OP' mode in which a low-level signal stops the output once the current sequence is finished and the last byte is read. We operate our AWGs in reset mode.

Unlike the prior hardware version, the two AWGs used in this setup do not have an incorporated bias-tee. To apply a voltage offset, an external bias-tee (Mini-Circuits ZX85-12G-S+) is used. This bias-tee was provided by the PPG manufacturer, after they made sure that the quality of the output pulses stayed unaltered. Since bias-tees act as high-pass filter, the concern was that long pulses would get strongly attenuated and distorted. For 50 ns long pulses no significant attenuation or distortion was observed.

The AWG can be programmed via a computer. Since the software provided by the manufacturer for the second-generation PPGs apparently is not compatible with our system, I wrote my own Matlab communication routine. The script creates the desired pattern, converts it to 8 bit hexadecimal and adds the necessary CRC (cyclic redundancy check) checksum before sending it to the device via serial over USB communication.

Electro-optic Amplitude Modulator Integrated EOMs are widespread in the field of optical telecommunications and in the last decade they have become available also for near-infrared (and visible) light. The advantages of integrated modulators are their compact size, the fiber-coupling, the fast response allowing for high modulation frequencies (up to GHz) and significantly lower voltage requirements than their bulk counterparts. They allow for modulation and creation of short optical pulses in the near-infrared.

The working principle of EOMs relies on the Pockels-effect, a second-order nonlinear process that describes the change of refractive index of a material when an external electric field is applied. Through a controlled change of the refractive index of the waveguide, the phase of the light traveling through it can be modified. Usually, this kind of modulator consists of a crystal sandwiched between two electrodes to which a voltage is applied. Lithium niobate ($LiNbO_3$ or LN) is the preferred material for optical-integrated modulators since it combines the necessary second-order nonlinearity with the material properties needed for waveguide fabrication.

Amplitude modulation can be achieved by inserting a phase modulator into the arm of a Mach-Zehnder interferometer. An applied voltage results in a phase difference between the two interferometer branches, leading to interference when recombining the arms, inducing a change in the optical output power. In order to turn the transmission of the amplitude modulator from *on* to *off*, or vice versa, the voltage V_π, corresponding to a phase difference of π between the two interferometer arms, is needed. Placing the electrodes of the modulator in both interferometer arms with opposite polarity, the so called *push-pull* configuration, allows halving the required voltage. The optical power after the modulator is described by the transfer function [255]

$$P = P_{min} + (P_{max} - P_{min}) \left(\frac{1}{2} \cos \left(\frac{\pi(V - V_0)}{V_\pi} \right) + \frac{1}{2} \right) . \tag{5.1}$$

Here V_0 is the closest voltage to zero for which maximal transmission is reached. The devices are fabricated trying to reach the maximum transmission at $V_0 = 0$, however, in practice deviations from this target value need to be controlled by applying a bias voltage. In order to generate pulses with such an amplitude modulator, a DC voltage of $V = V_0 \pm V_\pi$ needs to be applied for maximal suppression. At the desired point in time a voltage pulse is generated by the AWG and is mixed with the bias voltage. The applied voltage pulses are slightly distorted by the EOM transfer function, which can be compensated through the programmed shape if necessary.

The half-wave voltage V_π depends on the repetition rate of the applied pulse pattern. When adjusting the amplitude of the programmed voltage pulse, it is advisable to monitor the optical output while adjusting the maximum voltage. In agreement with the transfer function of the modulator, a double peaked optical output can be observed if the chosen voltage is too high. In this scenario, the maximum voltage creates a dip in the center of the pulse, since it is already on the falling edge of the transfer function.

In our quantum memory experiments we use fiber-coupled amplitude EOMs from Jenoptik (AM785) for the pulse-shaping. They reach a specified extinction ratio of around 30 dB. Due to the optical damage threshold the modulator tolerates a maximal CW optical power of 50 mW. A low optical input power is nevertheless preferable for a more stable operation of the device. The maximum RF voltages are ± 15 V.

Unfortunately, the V_0 of EOMs tends to drift. Several factors contribute to this drift, one of which is temperature fluctuations caused externally or internally, for example due to the RF signals and optical power heating up the device. These drifts can be partially compensated [256], for a reliable operation of the modulators on the time scales of interest, however, we opt for active feedback on the bias voltage.

Locking the EOM The operation point of an EOM can be stabilized through feedback on the bias voltage. In our quantum memory experiments the EOMs are used to generate the signal and control pulses, which are only a few ns long. Most of the time, corresponding to the state preparation which occurs before each storage attempt, the EOMs should maximally suppress the transmission. The generated pulses are short in comparison to the temporal duration of the experimental cycle which is in the order of µs. This low duty cycle allows for locking the EOMs to minimal transmission. For this purpose feedback reacting to slow drifts, averaging out the short pulses, is desired.

In this experimental setup a beam sampler picks off a small fraction of the output of the EOM which is then detected by an amplified PD with switchable gain (Thorlabs PDA36A2). The gain is chosen to best exploit the whole dynamic range of the PD without saturating it. To make the lock less sensitive to stray and ambient light a lens tube with a long-pass filter (cutoff wavelength 700 nm) is mounted on the PD. The PD signal is sent to the breakout box (NI BNC-2090A) of a computer integrated multifunction I/O card (NI PCIe-6363). A LabVIEW routine based on lock-in techniques processes the input and digitally determines the best bias voltage. To compensate the slow drifts of the EOM's working point kHz lock-in frequencies suffice.

Eventually, the bias voltage necessary to compensate the drift reaches the maximum rating of the modulator. In this case, the LabVIEW routine automatically performs a jump of $2V_\pi$ towards zero. We refer to these jumps as 'turnovers'. The timescale at which these turnovers occur depends on the incident power on the modulators. Attenuating the beams before the EOM reduces the frequency of the turnovers so

that they usually do not occur within the chosen integration times. The locking parameters are optimized by aiming at maximum suppression while at the same time making sure that the modulator automatically re-locks after a turnover.

The output of the I/O card is sent through an amplifier (aim TTi WA301), which is used solely for impedance matching. Subsequently, the bias voltage is combined on a bias-tee with the pulse generated by the AWG before reaching the EOM. The cable length between the devices is kept as short as possible in order to keep the response time of the overall pulse generation process short.

The presented technique for the generation of short optical pulses is applied to both the signal and control generation. Once the pulses are generated, they need to be attenuated to the single-photon level or amplified in order to fulfill their respective task.

5.2.3 Weak Coherent Pulse Generation

The topmost box labeled as 'signal' in Fig. 5.2 shows the experimental setup used for the generation and calibration to single-photon level of the signal pulses. For the generation of our signal pulses a DFB laser diode emitting in CW on the D_2 line is used. A beam sampler picks up a small fraction of the beam right after the laser for auxiliary purposes such as monitoring the emitted frequency on a wavelength meter (HighFinesse WS-7). A second pick-up leads to a Doppler-free saturation spectroscopy setup used as an absolute frequency reference when determining the vapor cell temperature or magnetic field strength through spectroscopy of the memory (see Section 4.4.3). The rest of the light is sent through an amplitude EOM which carves the programmed signal shape out of the CW light, as described in the previous section.

As a proxy signal, I chose to generate pulses that resemble the temporal shape of the signal photons emitted by the SPDC source as closely as possible. By doing so, this experiment represents a fair test-bed for later experiments conducted with actual single photons. I recorded the temporal shape of the optical pulses with the single-photon detector and compared it to the cross-correlation $g^{(2)}_{s,i}(\tau)$ of the SPDC source, which corresponds to the temporal profile of the heralded photons. It is important to bypass the signal filter stage for this measurement, since the etalons would broaden the pulse in the time domain. I programmed the electrical pulse as a bilateral exponential function and varied its time constants by trial and error until the measured optical signal reached a satisfactory similarity with the cross-correlation. By programming a pulse with an exponentially rising edge with a 1/4 time constant of $\tau^{\text{rising}}_{1/4} = 2\,\text{ns}$ and an exponentially decaying edge with $\tau^{\text{falling}}_{1/4} = 0.75\,\text{ns}$ a satisfactory similarity between the generated optical pulse and the targeted shape is achieved. Converted to more conventional quantities, as $1/e$ time constants, the values are $\tau^{\text{rising}} \approx 1.44\,\text{ns}$ and $\tau^{\text{falling}} \approx 0.54\,\text{ns}$ respectively. In Fig. 5.3 the programmed electronic pulse is compared to the pulse after applying the transfer function of the EOM as modeled by Eq. (5.1) and to the resulting temporal shape

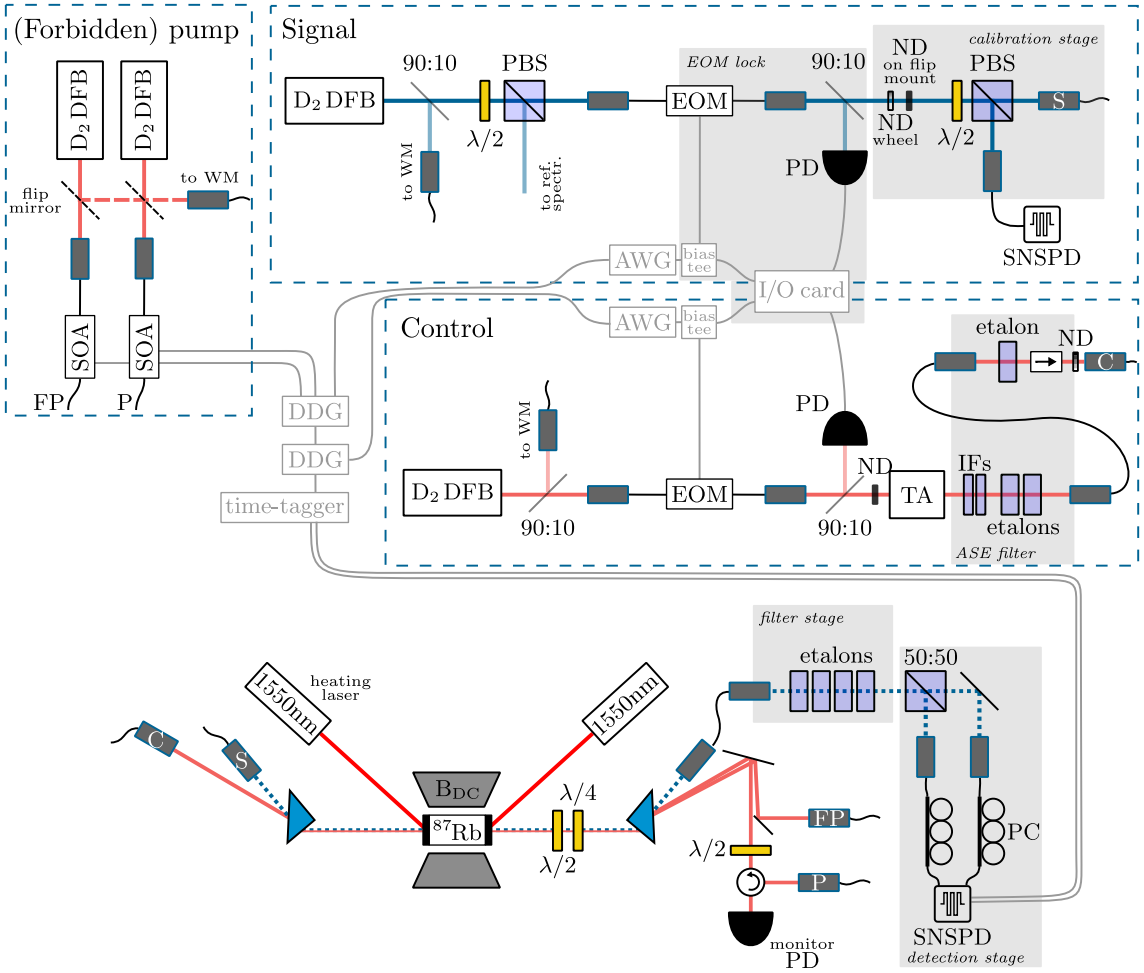

Figure 5.2: Sketch of the experimental setup of the HPB memory operated with WCPs. The dashed boxes in the top half show the preparation stages for the signal, the control, and the pump pulses. Detailed descriptions of these preparation stages can be found in the main text. The heart of the quantum memory setup is shown in the lower part of the figure. Light gray devices and connections represent the path of the logical trigger and follower pulses that control the memory. Abbreviations: PBS – polarizing beam splitter; WM – wavelength meter; ND – neutral density filter; EOM – electro-optical modulator; TA – tapered amplifier; IF – interference filter; PD – photodiode; SOA – semiconductor optical amplifier; DDG – digital delay generator; AWG – arbitrary waveform generator; X:Y – beam splitter with the specified branching ratio; PC – polarization controller; SNSPD – superconducting nanowire single-photon detector.

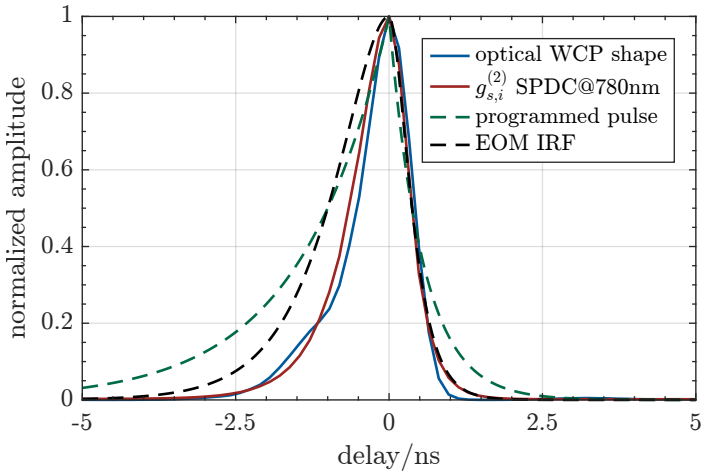

Figure 5.3: Comparison of the programmed voltage pulse with the temporal profile of the optical WCP used as signal proxy. A pulse with time constants 2 ns rise and 0.75 ns fall is programmed on the AWG (green). The pulse gets slightly distorted by the transfer function of the amplitude modulator (black). The resulting optical pulse (blue) is about 500 ps shorter than the programmed length. This offset appears to be a constant introduced by the AWG, independent of the pulse duration. By iteratively adapting the pulse parameters a good overlap with the cross-correlation of the SPDC source (red) is reached. The bin width of the histogram is 162 ps, with the time resolution given by the time-tagger.

of the optical signal. Furthermore, the figure includes the cross-correlation $g^{(2)}_{s,i}(\tau)$ used as a target for the signal shape.

For the operation of the memory with WCPs it is important to calibrate the mean photon number per pulse $|\alpha|^2$, to have reproducible experiments. In our experiments the generated pulses are attenuated in the calibration stage to contain on average one photon per pulse.

Calibration Stage After the amplitude modulation, the signal beam is sent to the calibration stage where it is split into two arms with a well defined power splitting ratio. The power ratio is set by rotating a half-wave plate in front of a PBS. Each arm is coupled into a PM fiber, one leading to the quantum memory and the other one directly to the single-photon detector. Usually, we choose a 1:1 ratio between the two arms. When calibrating the ratio of the two arms we measure the CW optical power after each of the fibers. This way we compensate for losses and differences in the coupling efficiencies between the two arms. Once the splitting ratio is set, ND filters are added before the PBS, attenuating both optical paths equally. Usually, we use a set of ND filters (nominal attenuation of OD 17.5) on a flip-mount for the substantial part of the attenuation, while the fine tuning is done with an ND gradient filter wheel (Thorlabs NDC-50C-4M-B, maximum nominal attenuation of OD 9.2). For knowing how much attenuation is required, it is necessary to estimate the expected detection rate for a given mean photon number. Since we work with

an attenuated laser, we generate WCPs. A coherent state can be represented in the Fock basis as

$$|\alpha\rangle = \exp\left(-\frac{|\alpha|^2}{2}\right)\sum_{n=0}^{\infty}\frac{\alpha^n}{\sqrt{n!}}|n\rangle \ . \tag{5.2}$$

The probability of having n photons in one pulse follows a Poissonian distribution described by

$$p(n) = |\langle n|\alpha\rangle|^2 = e^{-|\alpha|^2}\frac{|\alpha|^{2n}}{n!} \tag{5.3}$$

with $|\alpha|^2 = \langle n\rangle$. For a coherent state with $|\alpha|^2 = 1$ the fraction of pulses containing no photon is given to be $p(0) = 36.8\,\%$. Consequently, the probability of having at least one photon per pulse is given by $p_{\text{click}} = 1 - p(0) = 63.2\,\%$.

The expected detection rate can thus be estimated by the product $r_{\text{calib}} = f_{rep} \times \eta_{det} \times p_{\text{click}}$. Here f_{rep} is the repetition rate of the experiment, as programmed in the internally triggered DDG. For most of our storage and retrieval experiments $f_{rep} = 300\,\text{kHz}$. The detection efficiency η_{det} of the superconducting nanowire single-photon detector (SNSPD) is channel-specific. Channel 8 of the SNSPDs is dedicated to the calibration of the mean photon number per pulse and has a specified detection efficiency of $\eta_{det} = 85(3)\,\%$. With these parameters the targeted detection rate on the calibration detector channel is $r_{\text{calib}} \approx 1.61\times10^5\,\text{cts/s}$. If an arbitrary ratio rather than 1:1 were chosen, the expression above needs to be modified accordingly. The count rate on the calibration port is monitored on the detector while the attenuation is adjusted until the measured rate matches the target rate. Due to the limited extinction ratio of the EOM (the signal pulse constitutes only a third of the total counts) combined with the low experimental duty cycle of 2.5 ‰ for the programmed signal sequence, it is important for the calibration to only count the detections within a specific region of interest (usually 8 ns to approximately match the retrieval window we choose). Otherwise, the mean photon number is overestimated due to the great weight of the background. If the generated signal pulses are modified, e.g. their width or their shape, it is necessary to recalibrate the setup. In case the repetition rate of the experiment is changed, the number of expected detection rates has to be updated.

We checked the calibration on a daily basis before starting the experiment. It turned out to be stable over longer periods of time, especially the splitting ratio at the PBS. During the experiments, the timestamps of the calibration port are not recorded in order to keep the size of the recorded file manageable. A short, separate recording, containing exclusively the calibration data, is taken on measurement days.

Finally, the signal pulses are coupled into an AR coated PM fiber and brought to the quantum memory.

5.2.4 Control Pulse Amplification

The acceptance bandwidth of a quantum memory, in our scheme given by the width of the transparency window, needs to spectrally 'fit' the whole photon. If this

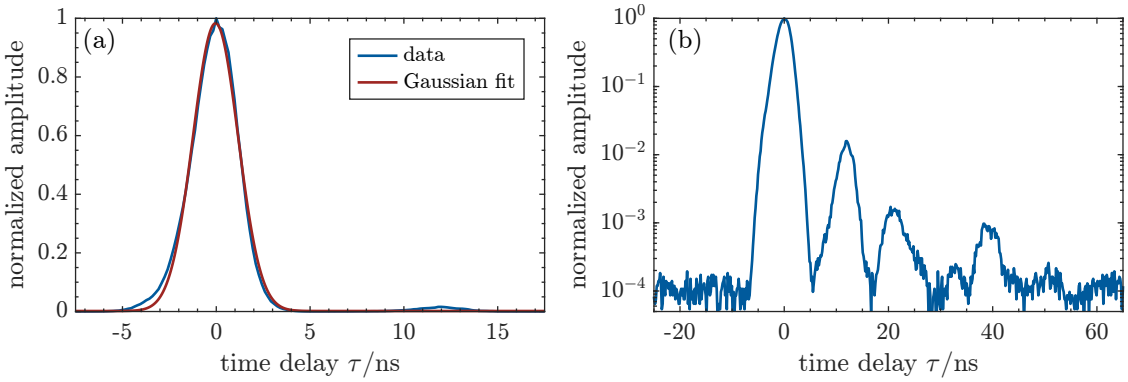

Figure 5.4: Shape of the amplified optical control pulses measured with the SNSPD after the first two etalons used for ASE filtration. A Gaussian pulse with a FWHM of 3 ns is programmed. **(a)** The measured data are fit with a Gaussian function. The fit yields a FWHM of 2.8 ns. **(b)** When plotted logarithmically the undesired ringing after the programmed pulse becomes obvious. The secondary peak is only a factor 100 lower than the principal peak – enough to perturb the spin wave.

matching is not given, only some spectral components of the photon can be stored and the process becomes inefficient. The width of the transparency window induced in the vapor, which strictly speaking is closer to the AT regime than EIT, increases linearly with the control Rabi frequency and thus scales as the square root of the optical power. Consequently, control pulses with high Rabi frequencies are required for broadband storage of single photons. When working with diode lasers, this most certainly means that the control field needs to be amplified. The experimental setup for the generation and amplification of the control pulses is sketched in Fig. 5.2 in the central box labeled 'control'.

For the generation of the control beam a DFB laser is used as well. Due to the optical damage threshold of the amplitude EOM of ~ 50 mW, determined by the fiber-waveguide interface, it is necessary to generate the optical pulses before amplification. The width and shape of the control pulses can be varied. We usually generate control pulses with a programmed Gaussian profile and a FWHM of 3 – 4 ns. The high temporal resolution of the DDG allows us to accurately time align the control pulse with the incoming signal photon. Figure 5.4(a) shows such a typical control pulse with a programmed FWHM of 3 ns and a Gaussian fit (red) yielding an actual FWHM of 2.8 ns.

The extinction ratio of the control pulse is important: If the control light leaks to the vapor cell in between storage and retrieval, the control can perturb the spin wave and lead to untimely read-out. Figure 5.4(b) shows the undesirable but typical dynamics of the control EOM. The actual programmed voltage pulsed is followed by considerable ringing.

In order to reach the high Rabi frequencies necessary to store the broadband photons in an EIT memory we rely on a tapered amplifier (TA). A TA is a semiconductor traveling-wave optical amplifier. Its gain medium is comparable to a laser diode, with

AR coatings on the end-facets of the semiconductor instead of partially reflective ones, in order to limit feedback – avoiding lasing. The gain region, which is tapered along one direction, is electrically pumped to achieve a population inversion. When seeded, the incident beam is amplified by stimulated emission. A divergent seed beam is coupled into the amplifier to match its tapered shape. In this way more charge carriers are available yielding a high gain and the intensity does not reach the damage threshold of the material [257].

The commercial TAs we use are intended for CW amplification by the manufacturer. In previous experiments [56, 29], however, seeding the amplifier with a pulsed input worked fine for us and reliably generated high-power pulses. Due to aging, the TA-chip of the optical amplifier (Toptica TApro 780 nm) I used for the measurements presented in the previous chapter started not to reach the specified CW output power anymore. However, after I replaced the chip of the TApro with one of the newer generation (TA-0780-3000-2), we started running into some issues when operating it with a pulsed seed. While the CW output of the new chip reached the specified $\sim 3\,\mathrm{W}$, the shape of the generated pulses deteriorated and appeared to randomly change with time (briefly blocking the seed input seemed to recover the original pulse shape, most of the time). Furthermore, a frequency dependence of the pulse shape was observed. In the troubleshooting process I tested the TA with different seed lasers, even with different types of lasers (DFB and ECDL). Furthermore, I ruled out the EOM and its driving electronics by using other devices and by monitoring the optical pulse shape before the amplifier. The problems persisted. As a last resort, in order to rule out other possible sources for this unexpected behavior, other than the optical amplifier itself, as a test we coupled the control pulses coming out of the EOM to the TA used for the Zeeman memory experiment (Toptica BoosTA pro 795 nm with TA-790-2000-1 chip). This test confirmed our suspicions: even if designed for a different wavelength, with this second TA the pulse amplification worked fine, thereby identifying the new TA chip as the culprit. Since these amplifiers are designed and tested for CW operation the manufacturer could not help us with a definitive diagnostic for the observed oddity. The two chip generations apparently have slightly different chip geometries, the newer 3000-series has a slightly larger tapering angle to allow for higher output power, which might lead to different behavior. Material-related effects could contribute and cannot easily be excluded[1]. However, these are all speculations that the manufacturer is not able to verify with their standard setups since they are not configured for ns-pulse operation. Since the BoosTA pro seems to do its job properly at 780 nm, even if its gain envelope is centered at 795 nm, further investigations about this matter were suspended, seeming beyond the scope of this work. With the Zeeman memory setup on halt because of the manuscript of [29] waiting for peer-review, this TA was available, so we decided to use it for the HPB memory.

When the EOM is set to full transmission, the TA can be seeded with $\sim 3\,\mathrm{mW}$ of CW light, resulting in an output power of $1.5 - 1.7\,\mathrm{W}$. According to the amplifier's specifications, the TA is not operating in the saturated regime at these input levels.

[1] private communication with manufacturer, mail, 07.10.2020

The output power is, however, high enough for our purposes. Furthermore, operating the TA in the linear regime has the advantage of the optical pulses getting less distorted.

In a TA lasing is suppressed. However, due to high gain, spontaneously emitted photons are amplified, resulting in a broad noise background known as ASE. As its gain region is nm broad (28 nm for the 6 dB below peak bandwidth), it is critical to spectrally clean the amplified control before combining it with the signal. ASE, in fact, contains light at the signal frequency. This means that a fraction of this noise is on resonance with the spectral filters after the memory and reaches the detectors. It thus cannot be removed afterwards. In order to suppress the ASE we spectrally filter the TA output with two 780 nm IFs (Laseroptik, 0.37 nm FWHM bandwidth specified by the manufacturer, angle tunable) and two 2 mm etalons (51(1) GHz FSR, 1150(20) MHz FWHM bandwidth). The IFs block a large portion of ASE but transmit both the control and signal wavelength. The etalons narrowly filter about the control frequency. The spatial mode of the TA output is cleaned by coupling it into an end-capped PM fiber. This end-cap consists of a short piece of fiber without core, which is spliced onto the end of a PM fiber. Without the core, the light is not confined and continues to converge. For optimal coupling the focus of the beam needs to be at the glass-glass interface, thus resulting in a much larger beam area at the air-glass interface compared to the mode field diameter of the bare fiber. The reduced power density allows the fiber to withstand higher coupled power levels without inducing damage. A third etalon is used to further filter the control pulses. To reach a good transmission this 8 mm thick filter (12.8(1) GHz FSR, 280(10) MHz FWHM bandwidth) needs to be placed after the fiber coupling, where the spatial mode of the beam is already cleaned up. Empirically we found that this last etalon plays a crucial role in spectrally cleaning the control pulses and reaching good SNRs in storage experiments, even though its bandwidth is comparable to the control's spectral width. This last etalon, however, also introduces crosstalk, which leads the SNSPD to 'latch' (see Sec. 5.2.8). This is caused by the counter-propagating pump beam, coupling in reverse into the control fiber and then being reflected by the plane reflecting surface of the etalon. Placing an OI after the last etalon prevents this effect. A second end-capped fiber brings the control pulses to the memory stage.

Due to the low duty cycle of the experimental sequence, the average incident power on the optics is significantly lower when the control is pulsed compared to CW operation. This large difference in incident power leads to differences in the optimal alignment due to thermal effects. The first alignment is done with a low power CW beam, chosen to be comparable to the average power of the pulsed operation. The transmission through the setup is then measured with a PD after each filter and fiber coupling and optimized in pulsed operation, if needed. Typical peak transmissions of the control pulses from the TA to the memory stage are in the range of 15 %. Losses are composed of 90 % and 80 % transmission through the first and second IF respectively, ~ 78 % per etalon and approximately 70 % per fiber coupling.

In order to determine the Rabi frequency of the control field we estimate the power

per control pulse. For a pattern composed of four control pulses, repeated at a rate of 300 kHz, an average power of $P_{\text{int}} = 550(10)$ µW is measured with a thermal power meter head (Thorlabs S401C) in the darkened lab, after subtracting the background. A total of $N_{\text{pulse}} = 1.2\times10^6$ pulses per second are generated. By recording the control pulse pattern with the SNSPD, and considering a $3 \times$ FWHM region per pulse, it can be estimated that $\eta = 5\,\%$ of the power lies outside of the pulses. The energy per pulse can thus be estimated by $E_{\text{pulse}} = (1-\eta) \times P_{\text{int}}/N_{\text{pulse}} = 435(8)$ pJ. For a Gaussian pulse with $\sigma = 1.2$ ns this yields a peak power of $P_{\max} = E_{\text{pulse}}/(\sqrt{2\pi}\sigma) = 145(2)$ mW. The peak Rabi frequency is given by $\Omega_{\max} = \frac{1}{\sqrt{2}\hbar}\sqrt{\frac{2I_{\max}}{c\epsilon_o}}d$, where the factor $1/\sqrt{2}$ comes from the fact that the control light is in a superposition of σ^+ and σ^- polarizations. The peak intensity can be expressed as $I_{\max} = 2\frac{P_{\max}}{\pi w_c^2}$, where the leading factor 2 comes from the fact that a $1/e^2$ waist is considered. For a control beam diameter of $2w_c = 440(9)$ µm and a dipole moment of the control transition $d = \sqrt{2}/3 \times 3.584 \times 10^{-29}$ C m the expression above yields a peak Rabi frequency of $\Omega_{\max} = 2\pi \times 683(15)$ MHz . This value constitutes the maximum peak Rabi frequency we can reach for the control for the given beam waist.

For changing the control Rabi frequency we use ND filters mounted on a revolving mount with discrete steps. This is more reproducible than the gradient ND filter wheel used before. The control pulses are attenuated prior to the last fiber coupling in order prevent that changing the attenuation affects the spatial overlap of signal and control beams. If substantial attenuation is needed, we move the ND filters upstream, before the EOM to reduce the number of turnovers.

5.2.5 Optical Pumping

From the early experiments in [56] it became clear, that using a dedicated pump beam instead of the control laser for state preparation is advantageous. First of all, in this way it is easily possible to choose different beam sizes for the pump and the control. While the control beam is focused in the center of the cell to yield high Rabi frequencies, it is preferable that the pump beam addresses a much larger volume. This way, the atoms that diffuse in from outside of the interaction volume during the storage process are also prepared in the ground state and won't generate noise during readout. Secondly, the control EOM is easier to lock with the resulting shorter duty cycle since it can be easily locked to minimum transmission. The state preparation takes times on the order of µs and is actually the limiting factor for the repetition rate of the storage experiments. With the aspiration to build a memory that needs as little changes as possible to be operated in a reactive scenario, fast (on the order of ns), binary optical switches are needed for the pump. To achieve this speed with concurrent high extinction ratio we use home-built semiconductor optical amplifiers (SOAs) based optical switches.

The pump on the forbidden transition is prepared and switched in an analogous way.

SOA Switches A SOA is basically an un-tapered TA. Their straight waveguide results in a lower gain, but at the same time a lower required operating current is needed (maximally 175 – 200 mA). It is composed of a semiconductor gain medium which is electrically pumped to populate the conduction band with charge carriers. When a seed laser is coupled into the semiconductor it gets amplified by stimulated emission. By cutting the applied current not only does the SOA not amplify the seed light anymore, it becomes absorptive. SOAs have already been used as fast optical switches for telecommunication applications at mid-infrared wavelengths [258].

This type of amplifier has a significantly smaller charge carrier density compared to a TA. Resulting in less dissipated heat, this fact allows for small and compact packaging as in a standard 14-pin butterfly package, without the need of external active cooling.

Since no commercial devices at the Rb wavelengths existed when we started, our in-house electronic workshop has designed and built a current driver for operating SOAs as fast optical switches. It can switch up to 200 mA in a few ns. It reacts to follower signals at TTL-levels and has no limit to the pulse length. In our applications, the current driver directly receives the logical pulses from the DDG. The fiber pig-tailed SOA (Superlum SOA-332-DBUT-PM) is directly mounted on a small heat sink on top of the driver. The optical damage threshold, with a maximum of 20 mW is rather low for an amplifier, and is easily reached if one is not careful when setting up the SOA. As usual in fiber-integrated devices, the weak spot is the fiber-semiconductor interface. To be on the safe side, I usually coupled only 300 – 500 µW into the SOA fiber and successively slowly increased the applied current, so that 20 mW are measured at the output. The current driver incorporates a Meerstetter TEC-1091 controller for temperature stabilizing the SOA through its integrated Peltier element.

Since the SOA becomes absorptive if the operating current is below a threshold value, the demands on the current extinction ratio are not as stringent as our requirements for the suppression of the optical signal. Used as optical switches the SOAs present an optical suppression of −50 dB to −60 dB at the storage time, corresponding to a few hundreds of ns after switching. Motivated by this behavior the switching of the SOAs is prioritized in the experimental sequence in order to guarantee the best possible suppression. A detailed description of the switching behavior of our SOAs with the home-built controller can be found in GB's thesis [254].

5.2.6 Vapor Cell as Quantum Memory

The heart of the quantum memory setup – the vapor cell, its heating technique and the electromagnet – remains the same as what I presented in Ch. 4 for the preliminary characterizations. The lower part of Fig. 5.2 shows the experimental setup around the vapor cell.

Signal and control beams are outcoupled and subsequently combined on a single right-triangle-shaped calcite prism (Bernhard Halle). Usually, these prisms are sold as cemented pairs (then called Wollaston prisms), with orthogonal crystal axes to compensate achromaticity. The induced displacement is negligible for the optical path lengths and the frequency differences of about 35 GHz involved in this scheme. As a reference: a displacement of $w_0/10$, where w_0 is the beam waist, was measured for a $\Delta\lambda = 15\,\mathrm{nm}$ at a distance of 30 cm. By using only a single prism, the strain-induced birefringence originating from the cementing of the two halves is avoided, resulting in a better polarization extinction. For the same reason, the prism is attached to its mount with a tiny piece of double-sided adhesive tape instead of being glued or clamped. A second such prism is used further downstream in the setup to separate the two beams again. Wave plates positioned after the vapor cell are used to compensate changes of the polarization in order to maximize the suppression of the control light after the second prism.

The calcite prism induces a significant astigmatism – an expected effect for a birefringent medium when omitting the second prism used for compensation. By placing the two prisms as closely as possible (about 250 mm on each side of the cell, as close as the magnet and the mid-infrared heating optics allow) around the vapor cell and choosing their relative angle carefully, to make the setup symmetrical, this effect can be compensated. With these precautions we were able to reach a fiber coupling efficiency into the signal arm of 85 %, compared to only $\sim 50\,\%$ that was reached when the prisms were farther away and their orientation not exactly symmetrical.

To reduce signal losses, the prisms are AR coated for 780 – 895 nm at an angle of incidence of 33 – 42° with a specified residual reflectivity of $< 0.5\,\%$ per facet. In order to reduce further losses, the number of optical elements in the signal arm is minimized by focusing the signal beam directly with the integrated fiber coupler (Schäfter + Kirchhoff 60FC-SF 18 mm) onto the center of the vapor cell. The waist of the beam is measured with a beam profiler (DataRay WinCamD-UCD12, pixel size $4.65 \times 4.65\,\mu\mathrm{m}^2$), accessible through a flip mirror, at an equivalent distance from the cell center. With the 2 mm short vapor cell used for this experiment it is harder to place the focus exactly on the longitudinal center of the cell, compared to a 7.5 cm long one. Out of convenience the control beam is also focused directly with the lens of the fiber outcoupler (Schäfter + Kirchhoff 60FC-SF 6.2 mm). A $1/e^2$ diameter of 185(9) µm (440(9) µm) is measured at the focus for the signal (control) at an equivalent distance of the vapor cell center. The diameter of the control beam is chosen to be decisively larger than the signal guided by the results of the numerical simulation I implemented (see Section 6.4).

A good overlap between signal and control beam can be achieved by aligning the calcite prism in reverse with diagonally polarized light. The two incoming beams are combined on the facet of the hypotenuse, while the overlapped beams exit through the facet corresponding to one cathetus. Before operating the memory, the overlap between the control and signal beams is checked and optimized on a daily basis. For this purpose the control beam, whose polarization is temporarily rotated by 90°, is coupled into the signal fiber leading to the filter stage. Since the modes of the signal

and control beam differ due to the different beam waists, the optical fiber acts as spatial filter due to the different degree of mode-matching, although the rejection is only about a factor of 3.

While the signal output is coupled into an AR coated PM fiber leading to the filter stage, a circulator is placed in the control path as described in the previous chapter. An optical isolator (Qioptiq FI-780-5SV) with removable covers for the 'blocked' beams is used as a circulator. The light coming from the memory in this mode is directed to a fast PD (Thorlabs DET02AFC, 1 GHz bandwidth), which is used to monitor the control pulses, while the pump beam, which has the same polarization as the control, is injected in reverse direction. The pump beam is aligned counter-propagating with respect to the control beam. The forbidden pump beam is overlapped with a D-shaped mirror with the pump under a small angle, trigonometrically estimated to about 2.6(9) mrad. The pump and forbidden pump beams are chosen to have a larger waist with respect to signal and control, so that atoms flying into the interaction volume during the retrieval are also prepared in the right state and do not contribute to read-out noise. The pump lasers are collimated to about 2 mm.

5.2.7 Filtration

According to the pulse energy determined above and $E = h\nu$ the applied control pulses contain $10^8 - 10^9$ photons. In order to be able to separate the retrieved single photon from the strong control pulse and the generated read-out noise, stringent filtration is necessary. So much filtration is needed, that the control pulse is blocked and cannot be detected. The noise should only be the one at the same frequency and polarization as the signal. Measuring the necessary suppression of the whole filtration stage at once with only one device is impossible with the available dynamic range. We therefore characterize the suppression of the single filtering elements one by one.

The filtration process starts right after the vapor cell. The aforementioned calcite prism discriminates the control beam by its polarization. A half-wave and a quarter-wave plate, mounted to have tilt degrees of freedom, are placed after the vapor cell in order to align the polarizations with the axes of the prism. By measuring the optical power after the signal fiber without the vapor cell in place, a maximum (minimum) transmission of 23.0(7) mW (0.80(2) nW) is obtained by rotating the wave plates, yielding an extinction ratio of $\sim 3.48(15) \times 10^{-8}$. A photodiode power sensor (Thorlabs S130C) was used for this measurement. For the maximum suppression measurement the lower end of the power meter's specified power range was almost reached. The best angle of the wave plates for achieving maximum suppression of the control needs to be regularly checked and adjusted. A flip mirror installed before the first etalon allows picking up a monitor signal for maximizing the polarization suppression.

In order to avoid additional dephasing of the spin wave, signal and control beam were chosen to propagate colinearly in this experiment. This configuration does not allowing for stringent spatial filtration as described in the previous section.

After the signal beam is coupled into a fiber, it reaches the filter stage, where we use cascaded etalons for spectral filtration. The idea of using monolithic etalons as mechanically stable spectral filters and to tune their resonances through thermal expansion is widely known [259, 260].

We use low finesse Fabry-Pérot interferometers constituted by a plane-parallel[2] silica substrate with a partially reflective coating on both sides with a reflectivity of 95(1) % specified at 795 nm for an angle of incidence of 0°. The etalons have a diameter of 15.0(1) mm. In order to tune the resonance frequency through thermal expansion we use a Peltier element. The TEC has a central through hole for the beam (UWE electronic UEPT-440-127-079E120) and is cemented between an aluminum mount containing the etalon and a copper (or aluminum) block acting as sink with thermally conducting epoxy. The temperature is stabilized with a Meerstetter TEC-1091 controller. The remaining fluctuations depend upon the operation temperature. Close to room temperature the etalons can be stabilized to about 5 – 10 mK. Up to temperatures of about 50 °C the fluctuations are acceptable. We use no enclosure to separate the etalons from the environment. Given the thermal expansion coefficient of silica a frequency shift of $-2.436\,\mathrm{GHz\,K^{-1}}$ is induced at 780.24 nm. With the given temperature stability the cavity resonances are stable to within 12 – 24 MHz. The thicker etalons can be tuned over several FSRs making it possible to always find a working point close to room temperature. For our thinner etalons we were lucky that the resonances could be tuned to the signal frequency with temperatures $< 25\,^{\circ}\mathrm{C}$. The etalons were fabricated for us on request by the optical workshop *Bernhard Halle*, while the metal components were manufactured by our in-house mechanical workshop. The flatness and parallelism is specified by the manufacturer to $< \lambda/20$. Given the measured finesse, this is on the conservative side. In fact, we need to assume much better surface quality to theoretically reproduce our measured data.

In the HPB memory setup the signal is spectrally filtered with four cascaded monolithic etalons of different thicknesses. Three 1.45(1) mm thick etalons are used. For these etalons a FWHM bandwidth of 1.19 GHz was measured, and each single one has a measured suppression of −33 dB at FSR/2. This particular thickness was chosen to result in an FSR matching two times the ground state splitting of the atoms in the high magnetic field, in order to maximize the suppression at the control frequency. The last etalon was chosen to be 4.00(5) mm thick (FSR = 25.5(4) GHz, FWHM bandwidth of 550(10) MHz) since its narrower bandwidth is closer to the spectral width of the signal photons, trimming the tails of the noise transmitted

[2] As the authors of [259] comment, the finesse of a plane-parallel etalon is limited by being on the edge of the stability condition diagram. Mode-matching can not be fulfilled on both mirrors, as this would take a plane wave. Choosing larger beam diameters can be a trade-off improving the mode-matching, but increases the surface quality requirements of the mirrors. The finesse is thus limited to $\mathcal{F} \approx 100$, which for our intentions is higher than what we aim for.

through the first etalons. Each single etalon reaches peak transmissions above 82 %, often even in the high eighties. The overall CW transmission through the filter stage of light at the signal frequency is around 38 % (including the fiber coupling to the stage and that to the detector), while the control is suppressed by approximately −200 dB. With this suppression no additional counts to the noise floor are detected when sending the control onto the unheated atoms.

For the spectral filtering to work it is crucial that the ASE background of the control pulse was already removed, since the ASE fraction at the signal wavelength would be transmitted by the etalon stage.

5.2.8 Detection

The memory output is measured in HBT configuration in order to have access to the arrival-time histogram as well as the second-order coherence of the detected photons. For this purpose a free-space 50:50 beam-splitter is placed after the last etalon. Each output of the splitter is coupled into a SM fiber with polarization controller connected to a SNSPD (Single Quantum EOS).

SNSPDs are state of the art single-photon detectors. Being commercialized only in the last decade, they combine high timing resolution, low noise, and high detection efficiency in an unrivaled way. They are composed of a thin film superconducting material, usually in the form of a meander line, that covers a large area for efficiently collecting the output of an optical fiber. The detectors are cooled well below their critical temperature and operated at a constant current, just below the critical current [261]. The absorption of a single photon locally breaks superconductivity, creating a 'hotspot' [262]. The current is redistributed around this disturbance and due to the narrow width of the nanowires the current density increases above the threshold value enlarging the non-superconducting region. Additionally, Joule heating leads to the growth of the resistive region [263]. This creates a resistive barrier across the nanowire. The current is redirected by the increasing resistance to the readout electronics, generating a measurable voltage pulse [261]. This redirection allows the nanowire to cool down, recovering superconductivity within short time (dead time $\leq$ 10 ns). If the incident photon rate is too high, the detector might not recover itself in the resistive state – the detector is latching. In this scenario the software resets the detector, resuming normal operation. A current review on SNSPDs can be found in [264].

Our detection system is composed of eight independent detectors within one cryostat. There are four channels optimized for timing resolution and four channels selected by the manufacturer for high efficiency. They reach timing jitters of $<$ 20 ps and $\leq$ 50 ps respectively. In our system the efficiencies measured by the manufacturer for the different types of channels do not vary significantly and reach values $>$ 80 %. The fast channels reach their low jitter by placing the electronic amplifiers in the 40 K stage of the cryostat unlike the rest of the driving electronics, which are at room temperature. The performance of the fast channels is especially interesting

#	Detection of	η_{SDE} at 785 nm [%]	Timing jitter [ps]	Dark counts [Hz]
2	Herald	84(3)	13(2)	< 10
5	η_h	82(3)	39(2)	< 10
6	HBT	91(3)	40(2)	< 10
7	HBT	87(3)	31(2)	< 10
8	Calibration stage	85(3)	50(2)	< 10

Table 5.1: Characteristics of the SNSPD channels used for the experiments presented in this chapter, as specified by the manufacturer. The jitter specifications have been verified experimentally.

for the herald photon detection, since a low jitter enables more precise temporal alignment of the control pulse with the signal photon. In fact the herald SPAD we used before was the limiting factor in terms of jitter budget. Compared to the SPADs used, these new detectors have a significantly lower timing jitter and not only higher but, crucially, more accurately specified detection efficiencies. The performance specified by the manufacturer for the detection channels used in this work is listed in Tab. 5.1. In the literature often the system detection efficiency (SDE) is used. It describes the overall efficiency, and includes the effects of the coupling into the device, the absorption by the detector (given by material and geometry) and the probability of the generation of an electrical output pulse after the absorption of a photon [261].

The detectors are enclosed in a closed-cycle cryostat and need to be cooled down to 2.7 K before operation. Before the cool-down process can be initiated, the cryostat needs to be evacuated. We use a turbomolecular pump (Agilent TPS-compact) to reach a pressure of $\sim 3 \times 10^{-5}$ mbar. A He-compressor (Sumitomo HC-4E) pumps highly pressurized He into the cryostat where it expands, cooling down the cold-head (Sumitomo RDK-101D). The low pressure He flows back to the compressor. It takes the system approximately 3.5 h to reach its operational temperature while the pressure further drops to 2×10^{-7} – 3×10^{-7} mbar. Once the operational temperature is reached, the turbo pump can be disconnected by closing a valve and then switched off. As long as there is no He leak in the system, the detector can technically be operated continuously until the cold-head or the compressor need to undergo maintenance due to mechanical wear (according to the manual, every 10 000 h and 30 000 h of operation respectively).

The voltage pulses are sent to an electronic driver by SMA cables. This electronic driver allows monitoring and controlling of the detection system through a web browser interface. It offers a real time monitor of the detected counts for each channel and displays the detector temperature. Furthermore, the bias current and the trigger threshold for each individual channel can be set. An automated bias current scan helps in finding the critical current and characterizing the behavior of

each channel. The electronic driver hosts the amplifiers for the channels optimized for detection efficiency and has an SMA output per channel. In our case, these outputs are connected either to the DDG trigger input or the time-tagger.

The time-to-digital converter (qutools quTAU) used for the memory experiments is still the same used with the SPADs in Ch. 3. This means we are not taking full advantage of the detector's time resolution since we are limited by the time-tagger's 162 ps bin size. For the heralded-photon operation, however, this should not affect the timing of the control pulse, and consequently the efficiency, since that will be triggered directly by the output pulse of the detector.

5.3 Storing Weak Coherent Pulses

For the storage of WCPs the pump and forbidden pump lasers shine on the vapor, preparing the initial atomic state. After a pumping time of the order of 2.8 µs the internally triggered DDG, which acts as the experiment clock in this scenario, sends a follower pulse to switch off the forbidden pump, and then 50 ns later the pump. This delayed switching ensures that the atoms moved by the forbidden pump into the $m_I = +\frac{3}{2}$ manifold do not remain in the state $m_J = -\frac{1}{2}$. These two lasers stay switched off for the whole duration of the follower pulse. In this time window the storage and retrieval experiment takes place, before the sequence is repeated. The DDG triggers the signal PPG to generate the WCP calibrated to $|\alpha|^2 = 1$. The trigger is relayed to the second DDG, which triggers the control pulse generation and finally triggers the time-tagger. The DDG output triggering the control AWG is operated in train mode, usually with two repetitions. For the measurements presented here the pattern programmed on the control AWG consisted of two identical pulses, 60 ns apart. This results in a total of four control pulses per sequence with a variable delay between the first and the last two pulses. The first pulse acts as a precursor and contributes to the optical pumping. The second pulse is time-aligned with the incoming signal pulse, acting as the read-in pulse. The AWG pattern is then repeated after a predefined train pulse spacing, which determines the storage time with the arrival of the third pulse functioning as read-out. The exact storage time can be set to a delay variable in increments of 20 ns, due to the functionality of the DDG's train mode, plus a small constant delay originating from the exact position of the peaks within the AWG record. For storage times shorter than the minimum spacing of 60 ns between the two train pulses, the write and read pulses need to be programmed in the same sequence of the control AWG. The last pulse is used as a reference for the noise level while monitoring the experiment with the live-view histogram during the initialization phase. It gives a rough idea of the level of the control-induced noise. Depending on whether or not this last pulse induces some read-out, conclusions about the retrieval efficiency can be drawn.

In the process of optimizing the operation of the quantum memory, a vast parameter space is involved. The characteristics of the control pulses can be tuned rather quickly, and to get a feel of their repercussions onto the memory performance, the

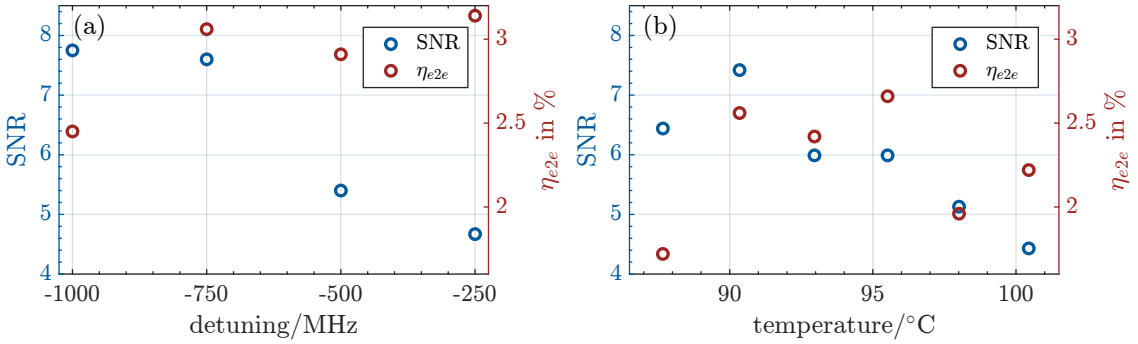

Figure 5.5: Sample sweeps of experimental parameters. **(a)** SNR and end-to-end efficiency for various working point detunings for an atomic temperature of 97(1) °C. For each measurement the control power and the time alignment of signal and control was optimized. The control pulse width was kept constant. The SNR appears to increase for larger detuning, while the efficiency becomes larger closer to resonance, as is expected. **(b)** SNR and end-to-end efficiency as a function of atomic temperature for a working point detuning $\Delta = -2\pi \times 750\,\mathrm{MHz}$. For each point the time alignment between signal and control pulse is optimized, while the Rabi frequency and the pulse width of the control are held constant. For the two highest temperature settings the forbidden pump laser was not used, since instead of improving the performance it increased the noise. These measurements were performed with an additional 8 mm etalon in the ASE filter stage, which was removed later on.

changes can be monitored on the live-histogram feature of the time-tagger. The control power is set by attenuating the beam with several ND filters on a revolving mount. By rotating the filter mount, different Rabi frequencies can be easily set. Once a suitable attenuation is found, it is convenient to place it before the amplitude modulator, thereby reducing the number of turnovers. The width of the control pulse can be swiftly changed by reprogramming the AWG. As the optimal pulse shapes obtained by a gradient ascent optimization in the scenarios studied in [178] were close enough to Gaussian pulses, an effort for further optimization was not deemed necessary, especially when considering the necessary time control. The width of the Gaussian control pulses had only little effect on the experiment, so that we changed it in steps of 0.5 ns. Decisively more important is the proper time alignment of the falling edge of the control pulse with the signal. For the initial temporal alignment, the control-induced noise was used as a proxy for the arrival time of the control pulses. The arrival time was shifted to roughly overlap with the incoming signal pulse by programming the necessary delay on the DDG. The fine tuning of the time alignment is successively performed while observing the height of the retrieval peak. When changing the optical setup, the temporal alignment needs to be updated accordingly: for example, adding an etalon in the ASE filter stage delays the control pulse by about 1 ns.

As we have already seen, the OD plays a crucial role in the efficiency of the storage process. In the short cell used here, the OD can be modified through the selected current of the heating lasers. A higher cell temperature will result in a higher OD,

which makes the process more efficient, but simultaneously causes undesired effects such as noise generation and increased collisions. A trade-off needs to be found. For this purpose I scanned the vapor temperature. Even though the cell is thin, we wait about 15 min after changing the temperature settings, in order to ensure that the vapor has thermalized. Figure. 5.5(b) shows a scan of the vapor temperature and the resulting SNRs and end-to-end efficiencies we measured. For this measurement series the working point detuning, the Rabi frequency and the pulse width of the control were held constant, while the time alignment between control and signal pulse is optimized for each point.

The working point detuning is a key parameter as well for the optimization of the performance. On resonance with the atomic transition, the storage and retrieval processes are expected to be most efficient. Concurrently, the noise contribution from the fluorescence should also be at its highest. In the absence of a simulation that models the noise processes, this parameter needs to be tuned empirically. In Fig. 5.5(a) the operating point is tuned, while keeping the vapor cell temperature and the control width constant. For each point the control power and the time alignment of signal and control pulses was reoptimized. The control pulse width was kept constant. The frequency detuning is the parameter in the setup that requires the longest time to be changed. Since we want to operate the memory on two-photon resonance, tuning the working point requires tuning all seven etalons involved in the memory setup. Even with the highly reproducible behavior of the etalons, and the capability to compute the necessary ΔT for a specific detuning, it remains the slowest parameter to sweep (and the largest effort).

Not only is the parameter space of the quantum memory huge, but the scan of a single parameter unfortunately gives only limited insight because of the interdependence of the parameters. For this reason, the goal of this optimization exercise was not so much to find an elusive global maximum, but rather to identify a good operating point to demonstrate our proof of concept. The simulation I implemented is supposed to help guide us in the future to find the right set of parameters more easily. For it to be complete and truly able to predict good operating points, a noise model should be implemented as well.

Apart from the optimization tasks mentioned above, the quantum memory setup requires regular maintenance. On a daily basis I realigned the overlap of the signal and the control beam and checked the transmission through the setup, e.g. the fiber couplings and the spectral filter stage. The laser frequencies were monitored throughout the measurements.

5.3.1 Storage and Retrieval

To characterize the performance of the memory I performed storage and retrieval experiments. The memory was operated at a detuning $\Delta = -2\pi \times 750\,\mathrm{MHz}$ from the excited state $|e\rangle$. With an operating current of 2.5 A per heating laser the atoms were heated to $T = 90(1)\,^{\circ}\mathrm{C}$. The weak coherent pulse input was calibrated

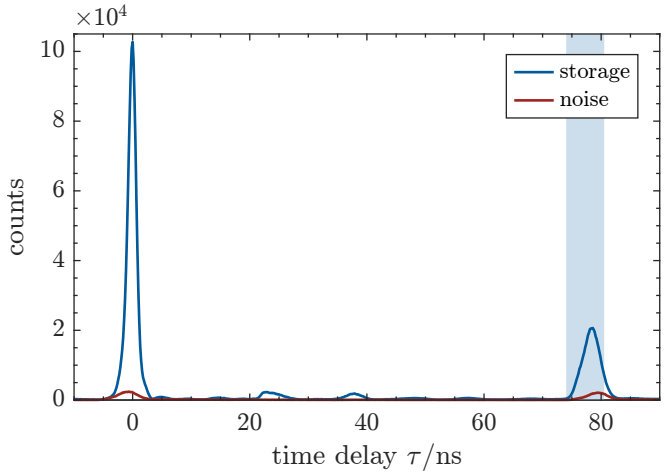

Figure 5.6: Arrival-time histogram of a storage and retrieval experiment with WCPs. The bin size is 162 ps. The signal pulse arrives at time $\tau = 0$ and is retrieved 80 ns later. The shaded area represents the 6.48 ns long retrieval region of interest. The histograms show the sum of the detections of the two SNSPD channels used in the HBT setup. The noise histogram is scaled to exactly match the trigger number of the storage and retrieval trace.

right before the storage and retrieval measurement. Within a region of interest of $100 \times 81\,\text{ps} = 8.1\,\text{ns}$ around the signal pulse a total of $N_{\text{calib}} = 4.80 \times 10^6$ counts is measured in the calibration arm. For an integration of 30 s this results in a rate of $1.59 \times 10^5\,\text{s}^{-1}$ yielding a mean photon number per pulse of $|\alpha|^2 = 0.97(6)$. For this measurement both pump and the forbidden pump on the forbidden transition were used to prepare the initial atomic state.

We directly measure the counts N detected by the SNSPD. To calculate the figures of merit we restrict our attention to a region of interest corresponding to the known time of attempted retrieval (a software-level time-gate). The detections of the two HBT detection channels are added. The noise measurements are corrected for the small variations in integration time with respect to the corresponding storage measurement by rescaling them.

In the measurement (see Fig. 5.6), the storage of a WCP for $\sim 80\,\text{ns}$ and successive retrieval was repeated at a rate of $f_{\text{rep}} = 300\,\text{kHz}$. Within an integration time of $t_{\text{int}} \approx 1\,\text{min}$ $N_{\text{trigger}} = 1.81 \times 10^7$ triggers and $N_{\text{ret}} = 4.46 \times 10^5$ counts were recorded within a 6.48 ns retrieval ROI. As an estimate of the read-out noise performance of the memory, we repeated the measurement and physically blocked the source. Within the same ROI $N_{\text{noise}} = 4.28 \times 10^4$ noise counts were detected. Using Eq. (2.9) this results in an SNR = 9.4(9). From the SNR the μ_1 parameter, describing the necessary mean photon number at the input of the memory, necessary to reach a SNR = 1 for the read-out [137, 58, 126]. The parameter is defined as $\mu_1 = \eta_{\text{HBT}}/\text{SNR} = 0.075(7)$ where η_{HBT} is the photon number expected to be detected for a WCP measured in HBT configuration (see end of section). With μ_1, the expected SNR for a single-photon source with a known outcoupling efficiency can be

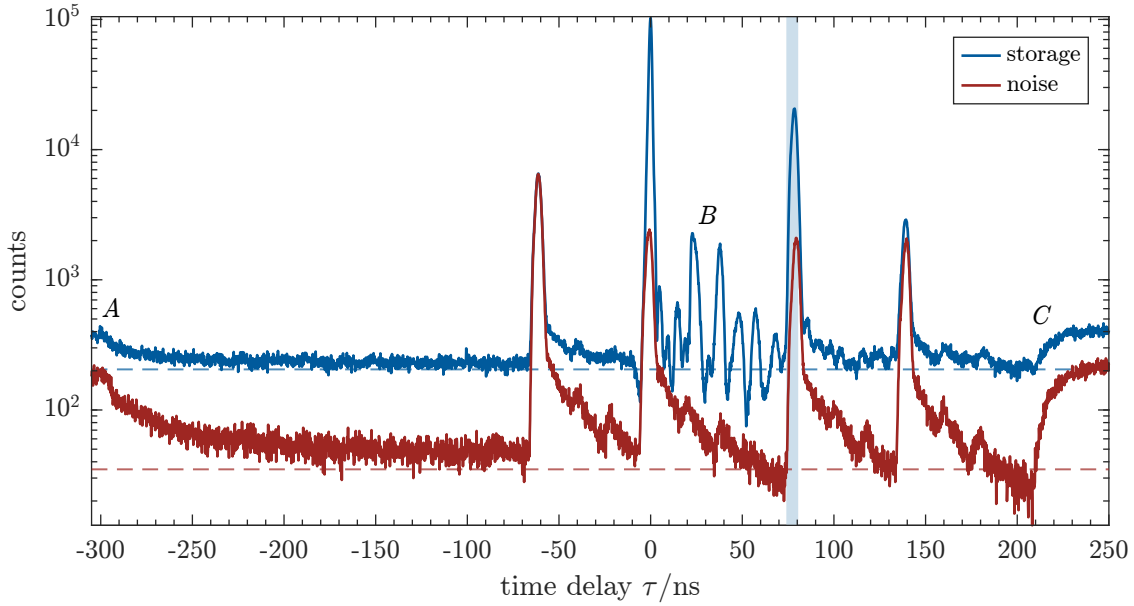

Figure 5.7: Arrival-time histogram of a storage and retrieval experiment with WCPs on a logarithmic scale. The bin size is 162 ps. The signal pulse arrives at time $\tau = 0$ and is retrieved 80 ns later. The shaded area represents the 6.48 ns long retrieval region of interest. The first and the last control pulse are reference pulses. The horizontal dashed lines are a guide to the eye for the maximum suppression reached by the EOMs. They show the average counts of the noise floor right before the first control pulse and right before the pump turns back on again. The offset between the two noise floors is caused by the limited suppression of the signal EOM. The labels A–C are described in the main text. The histograms show the sum of the detections of the two SNSPD channels used in the HBT setup. The noise histogram is scaled to exactly match the trigger number of the storage and retrieval trace.

estimated. This estimation is only valid if the input photons have similar properties and the memory is operated under the same conditions.

Figure 5.7 shows the corresponding arrival-time histograms recorded for the storage measurement and the noise trace on a logarithmic scale. Each plotted trace is the sum of the detections of the two channels used in the HBT setup. After 2.8 µs of state preparation, around $\tau = -300$ ns (label A), the pump laser is turned off for a fixed duration of 500 ns in which the storage experiment takes place. The forbidden pump laser is turned off 50 ns earlier. Its contribution to the noise was so little in this measurement, that it cannot be recognized in the histogram. After switching the pump lasers off we wait about 230 ns. This time, which is longer than the fluorescence from the state preparation needs to subside, is chosen to move towards limitations imposed by the future interfacing experiment with the SPDC source. At $\tau = -60$ ns the first control pulse is applied to aid the state preparation. The fact that it consistently induces significantly more noise than the following three control pulses suggests that it has a pumping effect. At $\tau = 0$ ns the input signal pulse arrives together with the time-aligned read-in control pulse. Around label B, during the storage time, unintentional retrieval can be observed. This retrieval is

induced by the ringing of the control EOM right after receiving the voltage pulse, which is translated in an optical signal. After the predetermined storage time has elapsed, the read-out control pulse retrieves the stored photon. By working with broadband single photons part of the fluorescence noise can be time gated due to the average excited state lifetime being longer than the retrieval time window . At $\tau = 140\,\text{ns}$ the reference control pulse arrives at the vapor cell. It can be seen that with the presence of an input photon at $\tau = 0\,\text{ns}$, more counts are detected at the time of the reference pulse. This could indicate that the shared atomic excitation is not completely mapped back into a photonic excitation during the read-out pulse and we observe some residual read-out at a later point in time. Around label *C* the pump and forbidden pump beams are switched back on and the experimental sequence restarts with the state preparation. The dashed lines represent guides to the eye that approximate the noise floor for the maximal average suppression achievable with the amplitude EOMs. By physically blocking the signal input this noise floor is significantly lower since only the light leaking through the control EOM remains, which is predominantly at the control frequency and is thus discriminated in the filtering stage. The noise measurement does not account for light leaking through the signal EOM when it should maximally suppress. I therefore correct the noise curve by adding the average offset between the two traces measured before the first control pulse. This offset contains 183 counts per bin, which results in 7325 counts within the considered region of interest. Adding this quantity to N_{noise} yields a SNR = 7.9(8) (corresponding to $\mu_1 = 0.09(1)$). A worst-case analysis can also be computed by using the reference (the fourth) pulse of the storage trace as a proxy for the noise. The fourth peak in the storage trace consists of the noise induced by the control and also contains the leakage of the signal EOM. However, since we are considering a time window 60 ns later than the retrieval time, the fluorescence induced by the control in the atoms might be stronger, since more time has passed from the state preparation. Furthermore, there might still be a small contribution of retrieval in this pulse, depending on how efficient the (intended) read-out was. This estimation can thus be regarded as a lower bound, resulting in a SNR = 5.9(5).

When measuring coherent states in an HBT setup with non-number-resolving detectors the detection probability changes compared to the scenario of a single detector. The 50 : 50 beam splitter of the HBT setup halves the intensity of the pulse. The new amplitude in one of the beam splitter outputs is therefore given by $\alpha' = \alpha\sqrt{\eta_{det}/2}$, where the losses due to the non unity efficiency of the detector are directly included. Since we sum the clicks of both SNSPD channels, the mean photon number expected to be detected per pulse when measuring with two detectors is thus given by

$$\eta_{\text{HBT}} = 2\left(1 - \exp(-|\alpha'|^2)\right) \,, \tag{5.4}$$

where the second terms follows from Eq. 5.3 as the probability for a pulse with α' to contain zero photons. Since the detector channels used for the HBT setup have slightly different detection efficiencies I approximate η_{det} as the weighted average of the channel efficiencies,

$$\eta_{det} \approx d\eta_{\text{CH6}} + (1-d)\eta_{\text{CH7}} \,, \tag{5.5}$$

where d is the ratio of the counts on the two detector channels. The efficiencies of the detector channels 6 and 7 can be found in Tab. 5.1. For the HBT setup this yields $\eta_{\mathrm{HBT}} = 70(4)\,\%$, for $d = 0.5145$, which is significantly larger than the detection probability on a single detector $\eta_{det} \times p_{\mathrm{click}} = 58(2)\,\%$ (for channel 6).

With Eq. (2.8) we can compute the end-to-end efficiency. From the detected counts an efficiency $\eta_{e2e}^{80\,\mathrm{ns}} = 3.18(18)\,\%$ is computed for the HPB memory in the retrieval time window. Correcting for the noise offset and the more conservative noise estimation result in only a slight change in efficiency, yielding $3.12(17)\,\%$ and $3.01(16)\,\%$ respectively. The readout induced by the fourth control pulse amounts to an efficiency of $\eta_{e2e}^{140\,\mathrm{ns}} = 0.13(1)\,\%$ (offset-corrected) for a 6.48 ns time window. Furthermore, the dynamics induced by the control EOM ringing during the storage time result in an involuntary readout leading to a loss in end-to-end efficiency of $0.45(5)\,\%$, accounting for some minor dynamics caused by the signal EOM.

5.3.2 Memory Lifetime

In order to determine the exponential $1/e$ lifetime of the quantum memory we recorded timestamps for storage and retrieval measurements with their corresponding noise measurements for different storage times. The memory was operated with the same parameters as in the measurement described above. The signal pulses were calibrated to $|\alpha|^2 = 0.97(6)$. The storage time is increased from measurement to measurement in increments of 40 ns. To ensure the same initial conditions for every storage time of the lifetime measurement we reduced the repetition rate to 100 kHz. This allowed us to program an off-time for the pump lasers of 1.5 µs – a lot longer than the times we expect to measure, so that the pump pattern stays unaltered for all chosen storage times. An integration time of 2 min was chosen for the measurements to compensate for the slower repetition rate. Figure 5.8 shows the measured end-to-end efficiency as a function of the storage time. The same offset correction described in the previous section is applied to achieve a more realistic estimation of the noise. The data is well described by an exponentially decaying fit yielding a lifetime $\tau = 224(8)$ ns. Dephasing of the spin wave can be ruled out as a limiting factor as a Gaussian decay would be expected. With a diffusion time of about 1 µs out of the beam, the lifetime is limited by the atomic motion out the interaction region.

Apparently, on the day of the lifetime measurement the setup was not optimized for efficiency equally well as it was for the day when the data for Fig. 5.7 was recorded. The efficiency for a storage time of 80 ns deviates by a factor 1.13 between the two data sets. This factor is used to correct the lifetime data in order to obtain a more realistic estimate of the internal efficiency (sometimes, like in the context of the simulation, also called total efficiency). By extrapolating the end-to-end efficiency to zero storage time a value of $\eta_{e2e}^{0\,\mathrm{ns}} = 4.7(3)\,\%$ is reached. After being freshly optimized, the transmission of CW light at the signal frequency through the whole signal arm of the setup (after the calibration point) is slightly below

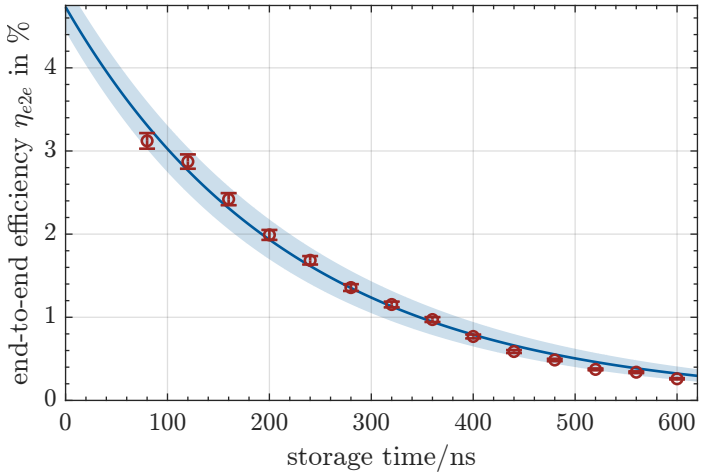

Figure 5.8: Lifetime measurement of the memory. The memory efficiency is measured as a function of storage time. The solid line represents a fit with an exponentially decaying curve with two free parameters (time constant and amplitude) that yields a lifetime of $\tau = 224(8)$ ns. The data is corrected by a multiplicative factor 1.13 to match the efficiency obtained for the 80 ns measurement presented above. Note that the lower efficiency obtained for the lifetime measurement, as compared to the 80 ns measurement, is probably due to a different degree of success in the optimization phase. The shaded area corresponds to the 95 % confidence interval of the fit. The intersection with the y-axis corresponds to the efficiency extrapolated to zero storage time.

20 %. Correcting for the transmission through the setup yields an internal efficiency for short storage times of $\eta_{\mathrm{int}}^{0\,\mathrm{ns}} = 24(3)\,\%$. According to the authors of [130], for the OD of this experiment the maximal internal efficiency achievable with forward retrieval is limited to $< 30\,\%$. The internal efficiency describes the efficiency of the storage and the retrieval process inside of the vapor cell, correcting for all the losses originating from the setup. This efficiency is the quantity to be compared with theoretical and simulated values. Note, however, that this is not a fair figure of merit for characterizing the memory in an applied scenario, since the memory cannot be operated without the surrounding setup.

5.3.3 Pump Characterization

The preparation of the initial state is essential for proper memory functionality. If the state $|s\rangle$ is still (partially) populated when the control pulse enters the cell, the corresponding atoms will be excited, leading to noise on the readout in form of fluorescence. With this in mind, we characterized the degree of the initial atomic polarization in the HPB memory with pump-probe measurements both in transmission and absorption.

For this measurement the signal laser was used as the probe. It is pulsed with the EOM into rectangular pulses approximately 50 ns long. In order to avoid a ringing

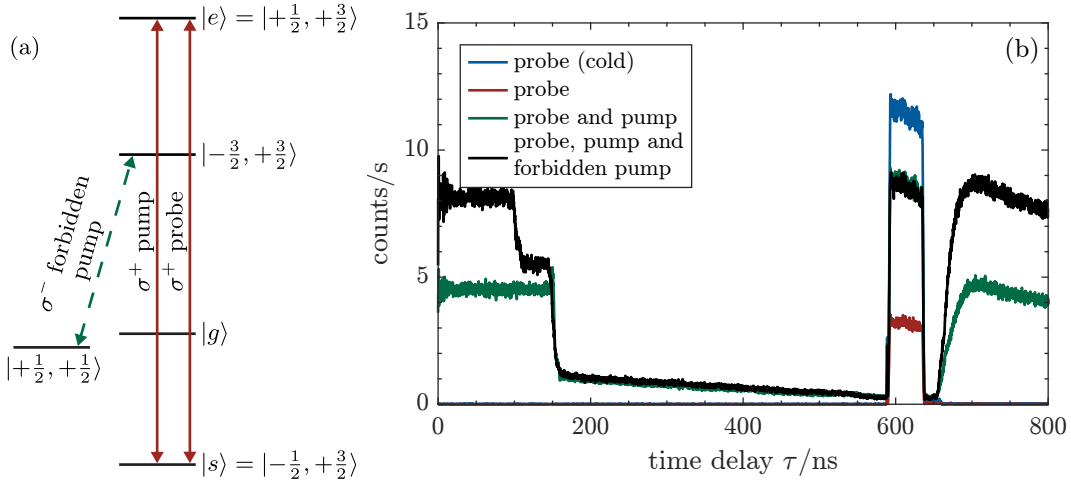

Figure 5.9: Characterization of the optical pumping in transmission. **(a)** Energy diagram of the involved atomic levels. The atoms are probed along the same transition to which the pump is resonant. The forbidden pump couples two different manifolds of the nuclear spin projection m_I through a forbidden transition. **(b)** Arrival-time histogram of the pump-probe measurement. The transmission of the probe in three different pumping scenarios is shown. The transmission of the probe through the unheated cell ('cold') is used as a reference. After switching off the pump laser, a decaying noise floor originating from spurious reflections of the fluorescence can be observed. This noise is specific to this characterization measurement and did not affect the storage and retrieval experiments. In the evaluation this background was subtracted from the data.

response that results in a spike at the beginning of the optical pulse, the rising flank of the voltage pulse needs to be programmed with a slight slope on the AWG. The frequency and polarization of the probe are chosen depending on the probing scheme – either along the pump or the signal transition. The pump and forbidden pump laser are operated as in the memory experiment, with the latter switching off 50 ns earlier. For these measurements the same repetition rate and the same pump dark time of 500 ns was chosen in order to have the same pumping conditions as in the storage and retrieval measurements. The probe beam is detected with the SNSPD after being appropriately attenuated to prevent latching.

Transmission With the first pump-probe experiment, performed in transmission, the depletion of the ground state $|s\rangle$ is assessed. For this purpose we probe along the pump transition. An auxiliary half wave-plate is added to the setup after the first calcite prism which rotates the probe's polarization in order to drive σ^+ transitions. In order to be able to detect the probe pulse without blocking the counter-propagating pump beam, the monitor port of the circulator is used to pick up the probe beam. In this scheme (see Fig. 5.9(a)) the more efficient the pump process, the higher the transmission of the probe should be. A weak probe is chosen. The probe pulse is attenuated to be still slightly transmitted through the unpumped hot vapor, resulting in a clear signal. Figure 5.9(b) shows the histograms of the different pump scenarios

considered. A reference measurement recording the probe alone, for which the vapor cell was not heated (referred to as 'cold'), is performed. For such a short vapor cell the OD is negligible at room temperature. This allows us to take a reference measurement without the need to physically remove the vapor cell. Subsequently, the atoms were heated up to 104(1) °C and the transmission of the probe through the unpumped ('probe') and the pumped ('probe and pump') atoms was measured. A further measurement was performed with the forbidden pump switched on as well. With this measurement scheme the beneficial effect of the forbidden pump should not be visible, since only the residual population in $|s\rangle$ is probed. However, it allows to investigate whether the forbidden pump has a negative effect on the atomic polarization of the $m_I = +\frac{3}{2}$ manifold.

By summing up the detections within a region of interest corresponding to the probe pulse, the OD of the probed transition and ultimately the atomic polarization for the $m_I = +\frac{3}{2}$ manifold can be computed. In panel (b) a decaying noise slope can be identified, which can be attributed to reflections in the setup[3]. In order to avoid overestimating the transmission, the noise floor is subtracted from the probe counts. The OD along the probe transition is given by $\mathrm{OD} = \ln(I_{\mathrm{in}}/I_{\mathrm{out}})$, where I_{in} is approximated as the counts within the probe pulse of the reference measurement. Likewise, I_{out} is the number of counts in the same time window for the pumped measurements. For the unpumped scenario an $\mathrm{OD}_{un} = 1.28(4)$ is measured, while the optical pumping reduces it to $\mathrm{OD}_p = 0.31(1)$. The forbidden pump only has a minor effect on the OD along the probe transition, resulting in $\mathrm{OD}_{fp} = 0.32(1)$.

As an assessment of the efficiency of the pumping process, the atomic polarization can be estimated from these results. The following consideration will be restricted only to the atoms in the $m_I = +\frac{3}{2}$ manifold. We assume that the two ground states are equally populated before pumping, due to a thermal distribution of the atoms. The OD of the unpumped probe transition can be related to the total number of atoms $N_{\mathrm{tot}} = N_{|g\rangle} + N_{|s\rangle}$ in the $m_I = +\frac{3}{2}$ manifold through $\mathrm{OD}_{un} \propto N_{\mathrm{tot}}/2$. Similarly, the OD of the pumped probe transition can be related to the residual atomic population in $|s\rangle$ through $\mathrm{OD}_p \propto N_{|s\rangle}$. The atomic polarization α for the $m_I = +\frac{3}{2}$ manifold can be thus expressed as

$$\alpha = \frac{N_{|g\rangle}}{N_{|g\rangle} + N_{|s\rangle}} = 1 - \frac{N_{|s\rangle}}{N_{\mathrm{tot}}} \approx 1 - \frac{\mathrm{OD}_p}{2 \times \mathrm{OD}_{un}} \,. \tag{5.6}$$

The measurements yield an atomic polarization of $\alpha = 87.9(5)\,\%$.

By repeating the same experimental procedure but varying the dark time between the pump and probe pulses, which is also known as Franzen sequence [168], the T_1 lifetime of the prepared state can be determined. In order to have the same pumping conditions for all dark times, a repetition rate of 10 kHz was chosen. The pump lasers were turned off for 20 µs and the dark time was swept.

[3] These reflections were exclusively present with the slight changes in the setup necessary for the pump-probe measurement. They were not observed in the storage and retrieval measurements.

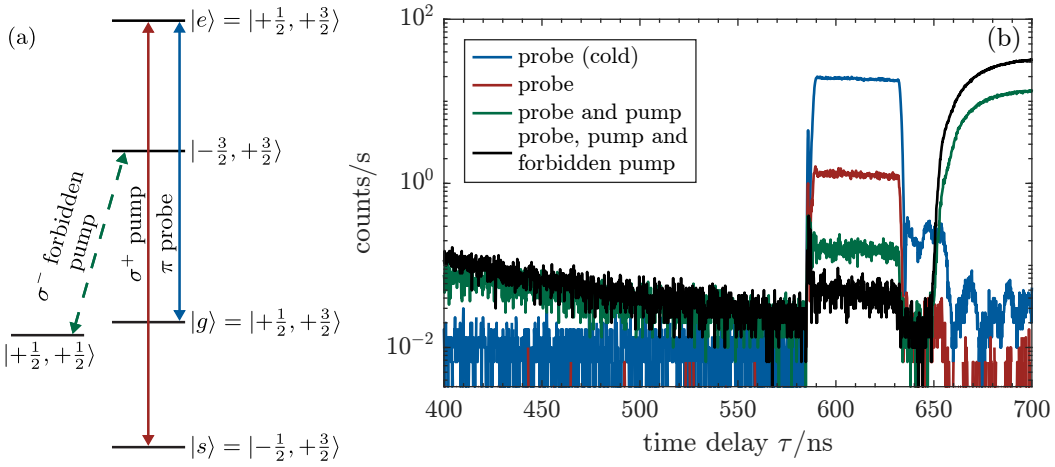

Figure 5.10: Characterization of the optical pumping in absorption. **(a)** Energy scheme of the involved atomic levels. The atoms are probed along the signal transition while they are pumped on the control transition. The forbidden pump couples the two different manifolds of the nuclear spin projection m_I through a forbidden transition. **(b)** Arrival-time histogram of the pump-probe measurement in absorption. The transmission of the probe in three different pumping scenarios is shown. The transmission of the probe through the unheated cell ('cold') is used as a reference. For the measurement with the pump and forbidden pump laser the signal is on the same order of magnitude as the noise. Theses measurements were not corrected for the background since slight differences in the choice of the average noise would have large repercussions on the computed OD. By not correcting the counts the OD is slightly underestimated.

For this measurement an atomic temperature of 104(1) °C was set and both pump and forbidden pump lasers were used. The data are well described by a single exponential fit with three free parameters (time constant, amplitude, and offset) yielding a lifetime of 6.9(6) µs. This behavior suggests that the lifetime is not limited by wall collisions, as close to the cell walls the relaxation process cannot be described by a single exponential decay [168]. The measured lifetime is shorter than what would be expected according to [168] by considering spin exchange collisions, wall collisions and buffer gas effects. This short lived state preparation, however, is compatible with measurements I performed at the beginning of my work, which lead to my decision to use another vapor cell for the imaging of the static magnetic field (cf. Appendix B). Nevertheless, since the the memory lifetime is more than an order of magnitude shorter than the lifetime of the state preparation, the latter is currently not a limiting factor.

In order to exclude the pump power at hand as the limiting factor for the achieved atomic polarization, I performed an additional transmission pump-probe experiment where two DFB lasers were tuned to the pump transition, each pumping with 20 mW. No significant difference in the measurement outcome was observed.

Absorption By choosing the probe to be resonant to the signal transition, on the other hand, the increase in OD induced by pump and forbidden pump fields can be analyzed. The pump-probe scheme is depicted in Fig. 5.10(a). No changes need to be applied to the memory setup, since the polarizations correspond to those of the storage experiments. Merely the filtering stage was bypassed.

The same three different pump scenarios as for the transmission measurement plus the reference signal were recorded. Figure 5.10(b) shows the various arrival-time histograms. The OD for the unpumped system is $\mathrm{OD}_{un} = 2.68(8)$. Due to the stronger transition strength of the signal transition compared to the control transition probed in the transmission measurement above, the higher OD is expected. For the pumped atoms ODs of $\mathrm{OD}_p = 4.74(14)$ and $\mathrm{OD}_{fp} = 5.90(17)$, including the effect of the forbidden pump, were measured. Under these conditions it appears that the forbidden pump laser increases the OD for the signal transition without deteriorating the atomic polarization.

As a consistency check, the result of transmission and absorption measurements can be compared. From the absorption measurement an unpumped OD of 2.68(8) resulted for the π signal transition. According to the transmission measurement, after pumping, the population in $|g\rangle$ is $2 \times \alpha = 1.76(1)$ times larger than in thermal equilibrium. The expected OD for the pumped vapor cell is thus $2\alpha \times \mathrm{OD}_{un} = 4.71(14)$, which is in very good agreement with measured OD in the absorption scheme.

An important phenomenon that needs to be considered when optically pumping optically thick media is radiation trapping [233]. In an optically thick medium, the photons emitted when the atoms relax from the excited state to the target state constitute a background radiation that can transfer the atoms back into the excited state and counteract the pumping effect. These fluorescence photons can be absorbed and re-emitted several times by the vapor. When the atomic density becomes too high, this multiple scattering competes with the pumping process and degrades the spin polarization that can be achieved. Imperfect state preparation then leads to Stokes noise during readout.

In [265] the effect of atomic and molecular buffer gases on the atomic polarization is studied. For atomic buffer gas species the achievable spin-polarization by increasing the buffer gas pressure saturates as soon as diffusion is no longer the dominant reason for the depolarization. Radiation trapping is then the limiting factor, and even increasing the buffer gas content further does not result in a higher polarization. Using a molecular buffer gas in the vapor cell, however, introduces quenching: a common technique used to limit radiation trapping by allowing the atoms to decay through non-radiative channels. It appears that in our vapor cell, filled with the atomic buffer gas Ar, we are limited by radiation trapping. To overcome this limitation a vapor cell filled with N_2 would be promising.

In [266] the authors studied the beneficial effects of quenching to avoid radiation trapping in the context of atomic vapor quantum memories. By using N_2 as buffer gas they were able to significantly push the highest atomic temperature, and thus

the OD, for which a near-perfect spin-polarization could be achieved. In Cs it is advantageous to pump on the D_1 line to further exploit the benefits of quenching, due to the higher chance of quenching collisions occurring with atoms in the $6\,^2P_{1/2}$ state. This can be explained with the larger quenching cross-section, the probability of quenching collisions occurring, for the D_1 line compared to the D_2, as well as the longer lifetime of the D_1 excited state making it more likely for a collision to take place when the atom is excited. In Rb, however, the quenching cross-sections for N_2 are equal within their uncertainties for the D_1 ($37(4)\,\text{Å}^2$) and D_2 ($36(4)\,\text{Å}^2$) line [267]. Furthermore, the difference between the excited state lifetime for the two D lines in Rb is smaller than in Cs. Form these parameters I expect less difference in quenching between the two Rb D lines compared to Cs.

The 'pancake'-geometry of the vapor cell used for these measurements is inconvenient in these regards. Usually, an elongated, thin cell would be preferred to have a high OD along the optical axis allowing the fluorescence induced by the optical pumping to easily escape in the transverse direction, thereby avoiding radiation trapping. For this proof of principle experiment we did not have the opportunity of freely choosing the cell design. For future developments, designing a suited cell geometry and an adequate buffer gas composition will be a key point.

An early attempt to use one of the SOA switches to generate the probe pulses, resulted in an unexpected behavior. A highly attenuated probe pulse seemed not to be affected at all by the optical pumping. Our conclusion was that the probe pulse was already so weak to be fully absorbed by the unpumped atom. The residual pulse we measured consists only of ASE generated by the SOA. With a width of about 30 nm, the ASE is sufficiently off-resonant not to interact with the atoms. Pumping the atoms would therefore show no effect on the ASE component of the pulse, assuming that the actual pulse is already fully absorbed by the atoms.

Effect of the Forbidden Pump In order to directly assess the effectiveness of the forbidden pump and the $m_I = +\frac{1}{2}$-pump laser on the memory functionality I performed a series of storage and retrieval measurements, all under the same conditions but successively adding pumping beams. The here presented data was recorded with an older version of the measurement setup compared to the one presented above and has mainly qualitative value. It was not fully optimized and acousto-optic modulators (AOMs) were used as optical switches instead of SOAs. The atomic temperature was set to 98(1) °C. Figure 5.11 shows the arrival-time histograms of the noise subtracted counts for the three measurements. The blue trace represents the scenario in which only the pump depleting $|s\rangle$ is on. The lowest retrieval peak and the largest leakage can be observed. By switching on the forbidden pump as well during the state preparation (red trace), less photons leak through the memory. These photons are not just incoherently absorbed. The retrieval peak is more than doubled in amplitude, evidencing that the storage process has become more efficient. By adding the $m_I = +\frac{1}{2}$-pump as well (see end of Ch. 4) the leakage peak decreases even further, albeit at the cost of a smaller retrieval peak.

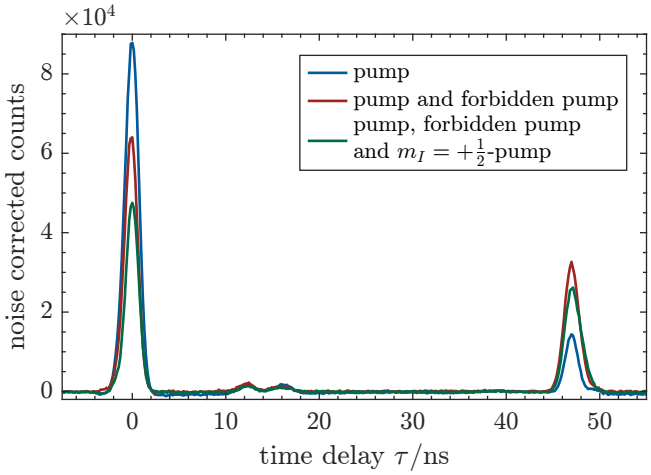

Figure 5.11: In an early storage and retrieval experiment the effect of the forbidden pump on the memory performance was assessed. While the forbidden pump lead to an increase of the noise-corrected retrieval peak, the $m_I = +\frac{1}{2}$-pump did slightly worsen the performance. The effect of the forbidden pump highly depends on the chosen parameters. The data were collected for 1 min.

The beneficial effect of the forbidden pump observed in this particular configuration, however, needs to be taken with caution. For higher temperatures, for example, this beneficial effect seems not only to vanish, but noise is added as well. Given the imperfect pumping within the $m_I = +\frac{3}{2}$ manifold of the ground state, adding atoms to it from the other nuclear spin states does not only contribute to a higher efficiency of the process, but can contribute to noise stemming from atoms in $|s\rangle$ as well. A trade-off needs to be made. Further investigation is needed to improve and understand the pumping scheme in order to take advantage of all the atoms in the vapor cell.

5.3.4 Noise Characterization

The tuning of the etalons we use as spectral filters makes them well suited for analyzing the different frequency components of a given light field. For a rather quick scan, the etalon can be ramped in temperature in a continuous way by setting the target temperature and enforcing a low maximum slope of the temperature ramp on the TEC controller. After discarding a few °C at the beginning and the end of the scan, the remaining part is linear. By always increasing the etalon temperature, hysteresis is avoided. The expanding etalon shows consistently reproducible behavior, allowing the recording of a separate reference measurement, where some light at the signal frequency is added, to calibrate the frequency axis of the actual scan. By scanning over at least one FSR in the linear region, the distance between the signal resonances can be used to convert the x-axis to a frequency axis. Figure 5.12(a) shows a noise spectrum recorded with this technique. In this scenario the first etalon (1.45 mm thick) of the filter stage is scanned. A flip mirror after the etalon allows us to bypass

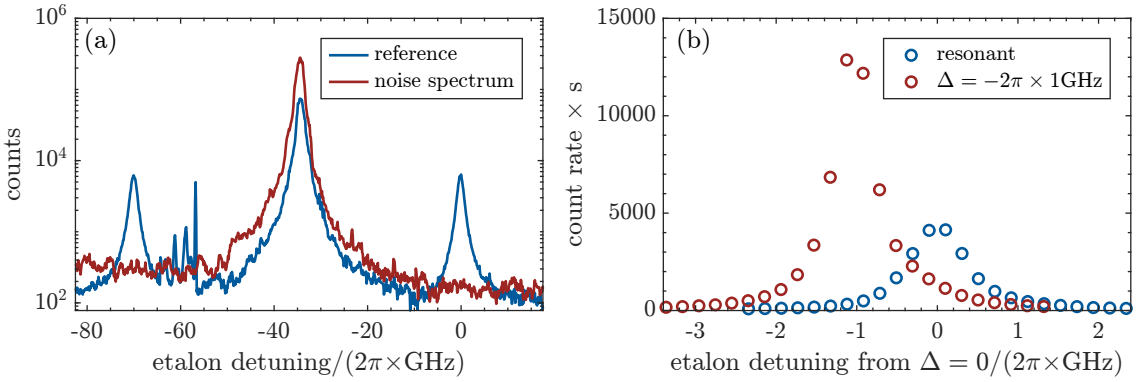

Figure 5.12: (a) Continuous scan of the first etalon (1.45 mm) of the signal filter stage. The noise spectrum is shown in red. The reference trace (blue), where some laser light at the signal frequency was added, was used to calibrate the x-axis through the known FSR of the etalon. The zero etalon detuning corresponds to the etalon temperature used for the storage experiments. The narrow features around $-2\pi \times 60\,\mathrm{GHz}$ in the reference trace can be attributed to turnovers of one of the amplitude EOMs. **(b)** Stepwise scan of the temperature of the last etalon of the signal filter stage. The measurement was taken with $\sim 2\pi \times 200\,\mathrm{MHz}$ steps, corresponding to 83.6 mK changes of the etalon temperature. Each single data point shows the time-gated noise counts integrated over the retrieval time window. The noise spectrum for two working point detunings is shown. The noise peak moves with the chosen detuning. On resonance the noise appears to be less.

the rest of the filters and to pick up the beam for detection. By scanning the first etalon, the polarization-filtered noise spectrum emitted by the memory is measured. The recorded spectrum is, however, broadened by the 1.19 GHz FWHM bandwidth of the filter. The etalon was scanned from 18 °C to 55 °C with a temperature ramp of $0.05\,^{\circ}\mathrm{C\,s^{-1}}$. The red trace shows the noise spectrum recorded for the memory working at a detuning $\Delta = -2\pi \times 1\,\mathrm{GHz}$ while the signal input was blocked. Apart from the main peak at FSR/2, corresponding to the control frequency, no other distinct features can be resolved with this broad IRF. For the reference trace (blue) some laser light at the signal frequency is added. The narrow features that can be seen in the reference measurement at around $-2\pi \times 60\,\mathrm{GHz}$ are an artifact caused by the turnover of one of the amplitude EOMs, since they are too narrow for being a real feature of the scan. In this type of scan the noise generated during the whole experimental sequence is integrated, including the fluorescence emitted during the pumping process. For the memory operation we are particularly interested in the composition of the noise during the retrieval window.

For this purpose, we have performed a time-gated noise measurement in which we scanned the etalon in steps. This is in order to distinguish between fluorescence caused by optical pumping and noise produced during retrieval, which is the only component of the integrated counts detrimental for operation. In this type of measurement both the TEC controller of the etalon as well as the time-tagger are interfaced with the computer and controlled via a Matlab script. The routine I wrote increases the etalon temperature one step at a time, comparing the setpoint with

the actual temperature, as read out by the TEC controller, to determine whether the etalon has thermalized. Once the temperature of the etalon is stable, the time-to-digital converter is triggered to start recording timestamps until a predetermined threshold value for the integration time (or file size) is met. The routine is looped to achieve the desired scan range. After the data are recorded, the detections are gated to the read-out window and are summed, constituting the noise level at one specific frequency point. This type of measurement requires more time to be recorded. Additionally, in the absence of frequency locked lasers, the drift of the lasers is monitored and compensated if necessary. Figure 5.12(b) shows a noise spectrum recorded with this second method. Each single point corresponds to 5 min of integration. In the presented plot the last etalon (4 mm) of the filter stage was scanned. The noise measured at this stage represents the residual spectrum passing through the whole filter stage except for the last etalon. The two traces represent two measurements performed with the memory operated on resonance and with a red detuning of $2\pi \times 1\,\mathrm{GHz}$. The x-axis is converted to frequency by using the temperature tuning coefficient, which depends on the thermal expansion coefficient of fused silica and is a function of the wavelength, and by setting the resonance frequency to zero. The recorded noise spectrum moves with the working point detuning. This characteristic is expected from stimulated Raman scattering and FWM. The noise peak is higher and narrower off resonance. It has, however, to be noted, that the experiments were performed on different days and that the experimental conditions might thus have been slightly different.

Both measurement techniques show the noise spectrum as modulo of the FSR. The measured noise spectrum can thus be complicated to interpret. For an extensive characterization of the noise spectrum, the same 'optical spectrum analyzer' setup I discussed in Ch. 3, composed of cascaded etalons, could be used. This technique is also of particular interest for analyzing the noise spectrum of the HPB memory, once a higher spin-polarization is achieved with the state preparation.

5.4 Interfacing with SPDC Photons

After obtaining the first evidence that we were able to store WCPs attenuated to the single photon level in such a short vapor cell with a high SNR, we decided to interface the memory with the photon-pair source.

For interfacing the quantum memory with the heralded photon source, the main DDG is programmed to accept an external trigger: in this way the detection of a herald photon starts the storage sequence. Whereas with a deterministic source, such as an attenuated laser, the experiment can be repeated periodically, for a probabilistic source, like SPDC, the memory needs to react to the detection of a herald. As a direct consequence, there is no earlier point in time, prior to the herald detection, where a device with a long insertion delay could have been triggered. For a periodic trigger, on the other hand, delays can be compensated by generating the

trigger earlier. In this way, even delays longer than the experimental sequence can be compensated.

For a probabilistic source this implies that the herald photon needs to arrive at the detector before the signal photon reaches the memory. Furthermore, the delay between herald detection and arrival of the signal should leave enough time for the various devices to react. On the side of the herald photon, the time between generation and detection can be kept as short as possible by building a compact setup and by choosing fibers as short as possible. This is usually only a small contribution. To guarantee that the memory, including all the electronics, has enough time to react, the signal photon is additionally delayed. A delay line in the form of an optical fiber grants a delay of about $5\,\mathrm{ns\,m^{-1}}$. For simplicity, we used the same 60 m fiber with AR end-facets as in the interfacing experiment in [29], giving us a measured delay of 285 ns.

A delay just short enough to allow the electronics and pulse generation to switch is desirable for the following reasons: Optical fibers have non negligible losses at NIR wavelengths of about $4\,\mathrm{dB\,km^{-1}}$. Delays on the order of hundreds of m already have considerable attenuation (only 91 % transmission for 100 m of fiber), which worsen the heralding efficiency of the source. Furthermore, for measurements between subsequently stored and retrieved photons, such as a HOM interference measurement to determine their indistinguishability, the delay of the signal photons sets a requirement on the memory lifetime and the repetition rate of the source. In fact, the delay should not be longer than the memory lifetime. It is also worth mentioning that longer delays also reduce the maximally achievable repetition rate of the experiment.

When selecting the electronic devices and the optical switches presented at the beginning of the current chapter, great care was taken to ensure that the memory logic would be able to react to a spontaneous trigger. In particular, this implied that the insertion delays of the instruments and the switching time of the optical switches were chosen to be faster than the adopted signal delay. Almost no modifications to the setup had to be made for the switch to SPDC photons, since this memory was planned from the beginning for a probabilistic source and with the experience from Zeeman memory in mind.

The temporal alignment of the signal pulse and the control pulse also plays a role in the optimal working point of the memory. Jitter in the electronics responsible for generating the control pulse leads to a fluctuation of this time alignment. In order to guarantee the stability of the temporal overlap of the falling control edge with the signal pulse, a low-jitter detector channel is used for herald detection. By doing so the largest jitter contribution for the control pulse generation is eliminated (compared to the operation with a SPAD). The current limitation should now be the AWG. For a detailed jitter budgeting of the devices used in our memory setups, I refer the reader to Section 5.2.1 of [254].

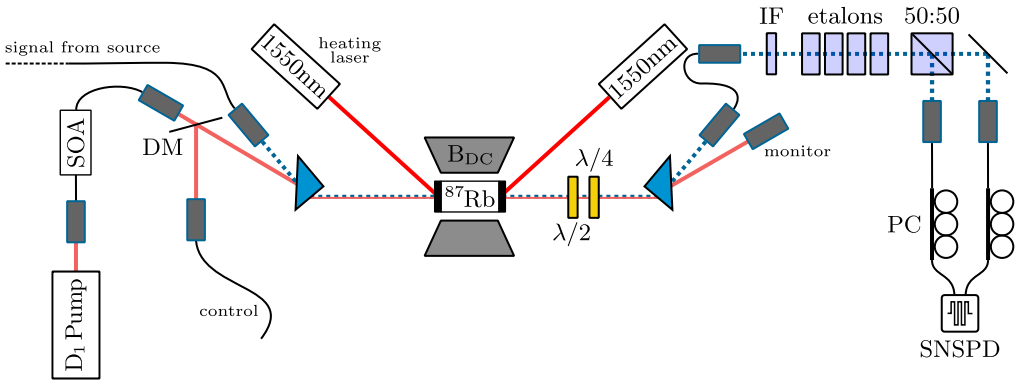

Figure 5.13: Experimental setup of the memory when interfaced with the SPDC source. The setup of the photon-pair source is described in Section 3.4 and is illustrated in Fig. 3.12. The detailed preparation of the control pulses and their ASE filtration is the same as in Fig. 5.2. Upon the detection of a herald the memory is triggered: the pump laser and the SPDC source are switched off and the control pulse is generated. The detection in HBT configuration allows for measuring the second-order correlations of the memory readout. Used abbreviations: DM – dichroic mirror; SOA – semiconductor optical amplifier; IF – interference filter; 50:50 – beam splitter with the specified branching ratio; PC – polarization controller; SNSPD – superconducting nanowire single-photon detector.

5.4.1 Pump Crosstalk

Both the pump as well as the forbidden pump induce crosstalk with the SPDC source. This manifests itself as a doubling of the herald rate compared to the scenario where these lasers are physically blocked, even though there already is an OI in the signal arm. By adding a second OI in the signal arm we can suppress this effect. However, we only reach ∼85 % transmission, which is detrimental for the heralding efficiency. Instead we preferred to change the pumping scheme. By using a pump laser on the ^{87}Rb D_1 line, we can easily switch to a co-propagating pump configuration while maintaining the same amount of available pump power, resulting in practically the same degree of atomic state preparation. For pumping on the D_1 line, a home-built ECDL with AR coated laser diode (Toptica LD-790-0120-AR-2) was used. The D_1 pump can be easily overlapped onto the same spatial mode as the control without significant loss of optical power by using a dichroic mirror (Semrock RazorEdge LPD02-785RU-25 with specified transmissions $T_{D1} = 97.4\,\%$, $T_{D2} = 0.52\,\%$). By combining the control and the pump beam after the fiber coupling leading to the experimental stage, the waist in the vapor cell of the respective beam can be individually set by choosing the focal length of the respective outcoupler. During the first storage attempts the forbidden pump laser appeared to have negative repercussions on the SNR. For this reason the forbidden pump was not used for the following measurements. An interference filter was added to the spectral filtering stage, since the pump induced noise leaked through the etalons and caused the detectors to latch.

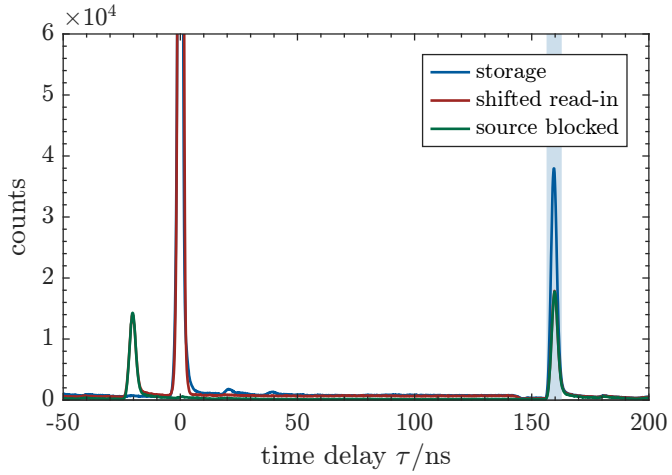

Figure 5.14: Arrival-time histogram of a storage and retrieval measurement of the HPB memory interfaced with the SPDC source. The bin size is 162 ps. Two noise measurements are performed: for the red trace the read-in pulse is shifted to arrive 20 ns earlier than the signal, while the green trace shows a measurement for which the SPDC source was physically blocked. The (cut-off) leakage peak reaches a peak value of 2.52×10^5 counts in the storage scenario, while 3.81×10^5 counts are measured when the read-in pulse is shifted.

5.5 Storing Heralded Single Photons

In the idle state of the memory, when operated in the reactive mode, the pump laser is on and continuously pumps the atoms into the initial state. Upon detection of a herald, the DDG first relays the trigger to the Pockels cell and the SOA switches. The switching of both the SPDC source and the pump laser is prioritized because of the insertion delay and the increasing suppression with time respectively. Subsequently, the control pulse generation is triggered and finally the time-tagger records the herald timestamp.

For a storage and retrieval measurement, the photon-pair source was operated with a pumping power of 4.5 mW, resulting in a heralding rate of $2.26 \times 10^5\,\mathrm{s}^{-1}$. The same working point detuning as for the WCP storage experiments of $\Delta = -2\pi \times 750\,\mathrm{MHz}$ was used. The atomic temperature was empirically optimized to $T = 93(1)\,^\circ\mathrm{C}$. A storage time of about 160 ns was chosen, as it is the minimum storage time when switching the uncorrelated noise floor. The corresponding arrival-time histogram is shown in Fig. 5.14.

For an integration time $t_{\mathrm{int}} = 8\,\mathrm{min}$, a total of $N_{\mathrm{herald}} = 1.09 \times 10^8$ herald detections were recorded, while the sum of the counts within the read-out window of 6.48 ns was $N_{\mathrm{ret}} = 7.23 \times 10^5$.

The noise of the interfaced system was measured in two ways. The first way to characterize the setup noise, probably the most widespread, is to physically block

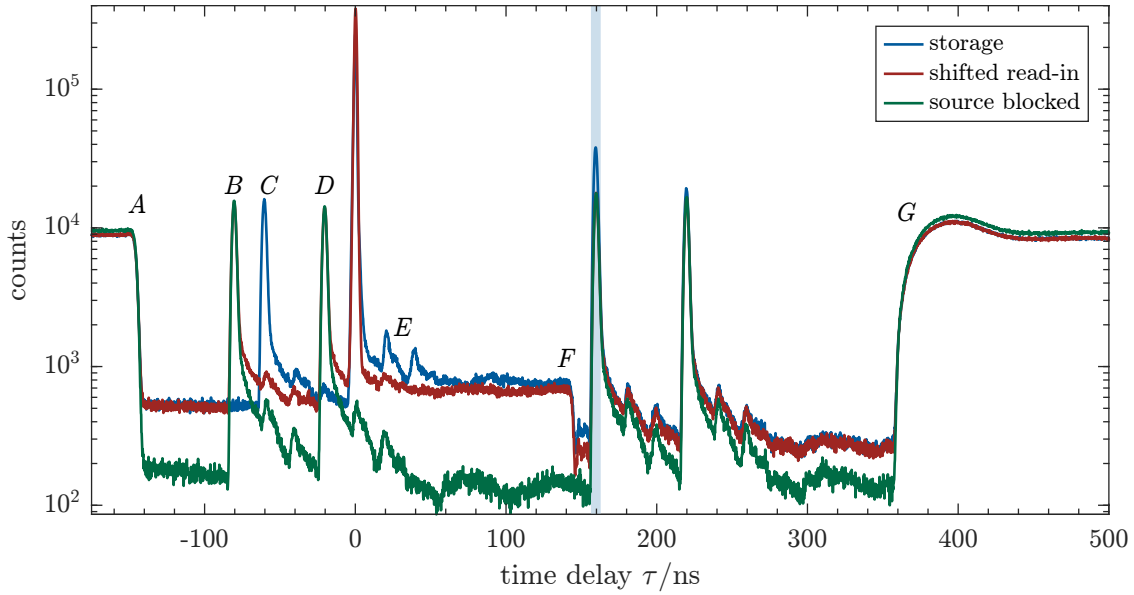

Figure 5.15: Arrival-time histogram of a storage and retrieval measurement of the HPB memory interfaced with the SPDC source on a logarithmic scale. The bin size is 162 ps. Two noise measurements are performed: for the red trace the read-in pulse is shifted to arrive 20 ns earlier than the signal, while the green trace shows a measurement for which the SPDC source was physically blocked. Within the retrieval time window the two noise measurements are very similar. Labels *A*–*G* are described in the text.

the heralded photon source[4], as we did in the WCP measurement. Within the read-out time window a total of $N_{\mathrm{n,block}} = 3.47 \times 10^5$ noise counts were recorded, resulting in a $\mathrm{SNR}_{\mathrm{mem}} = 1.08(11)$. By performing the noise measurement this way, however, the noise could be underestimated. In fact, the interface between source and memory is prone to crosstalk between the two systems. This crosstalk is eliminated by physically blocking the signal path and only the noise originating from the memory itself would be accounted for. To capture the noise of the interfaced systems, a second noise measurement is performed. Therein the read-in pulse (and the corresponding reference pulse, for technical reasons) is shifted to arrive 20 ns earlier. Due to this temporal misalignment, the single photons travel through the memory without being stored: they are either incoherently absorbed or transmitted due to the detuned working point (see the leakage peak in the noise trace at $\tau = 0$ ns). Since the atomic response to the read-out pulse is also influenced by the write pulse, we decided to shift the read-in pulse instead of omitting it. This keeps this systematic error as small as possible. For such a measurement with a shifted read-in pulse, $N_{\mathrm{n,shift}} = 3.52 \times 10^5$ noise photons were detected within the retrieval window, yielding an $\mathrm{SNR}_{\mathrm{tot}} = 1.05(11)$. There is almost no difference between the two noise characterization measurements.

[4]To be more specific: Since the noise measurement is also triggered by the detection of a herald, what is actually blocked is the signal arm, after signal and idler photons are separated by polarization.

In the histogram shown in Fig. 5.15, $\tau = 0\,\text{ns}$ is chosen to coincide with the arrival time of the signal photon. The steep edge that can be seen below label A is due to the pump laser switching off. The programmed off time is 500 ns long. As in the WCP scenario, the applied control pattern consists of four pulses. B and C mark the precursor pulse and its shifted counterpart in the noise measurements respectively. D is the time-shifted read-in pulse for the noise characterization. At zero time delay, we see leakage of the signal photons in both the blue and the red trace. After the read-in pulse, the features underneath E represent some unintentional readout. As in the WCP scenario, the ringing of the control EOM induces some retrieval right after the storage. At the time labeled F, the SPDC is turned off by the polarization rotation of the downconversion pump by the Pockels cell. The insertion delay of the Pockels cell is currently the element dictating the minimal storage time, while suppressing the uncorrelated noise floor, due to its insertion delay. The shaded region represents the 6.48 ns retrieval time window in which the read-out pulse arrives. The last pulse is a reference pulse. The storage curve appears to have slightly more counts than the noise curves, which could be explained by the occurrence of some retrieval. The difference, however, is too small to enable a definitive statement. Label G signalizes the turning on of the pump laser. The Pockels cell switches back on outside of the region shown in the figure. While within the region of interest both the noise curves have similar numbers of counts, it can be noticed that when the Pockels cell switches off the source, the noise floor is higher than when the source is blocked.

For computing an accurate end-to-end efficiency, the heralding efficiency of the source was measured right after the storage measurements. In a $80 \times 81\,\text{ps}$ long coincidence window $\eta_h = 42.8(16)\,\%$ was measured. Considering that this efficiency was measured with the delay line and the OI in the signal arm, it is as good as could reasonably be expected. Detector channel 5 was used for this measurement and 9.34×10^6 coincidences were measured for 2.66×10^7 detected heralds. Using the measured η_h and the average of the two HBT detectors' efficiency weighted by the detector ratio as η_{det}, Eq. (2.8) yields an end-to-end efficiency of $\eta_{e2e}^{160\,\text{ns}} = 0.90(5)\,\%$.

In order to characterize the state accuracy of the retrieved photons, the conditioned second-order autocorrelation is computed. A $g^{(2)}_{i:s,s} = 1.290(33)$ is obtained for a single 6.48 ns bin. The photon number statistics are dominated by the read-out noise. A total of 1542 triple coincidences were detected in the integration time. From the noise characterization with the time shifted read-in pulse a $g^{(2)}_{\text{noise}} = 1.81(8)$ for the noise is obtained. The expected value for the $g^{(2)}$ for an incoherent noise mixture can be estimated by using Eq. (2.10) where $g^{(2)}_{\text{input}} = 0.0298(1)$ was taken from the SPDC source characterization. Using the equation mentioned above a $g^{(2)}_{\text{theo}} = 0.94(2)$ is expected, assuming a purely incoherent background. This strong underestimation of the $g^{(2)}$ indicates that coherent noise sources are present and a model that accounts for coherent processes, as in [55], should be used for comparison to the data.

5.6 Discussion

The presented proof-of-principle implementation of a quantum memory in the HPB regime has demonstrated suitability for WCP storage in a short vapor cell. At the beginning of the project it was not clear whether a micro-fabricated cell would allow for storage and retrieval. By engineering a near-ideal three-level system, working in a high bandwidth regime where the signal pulses are temporally significantly shorter than the excited state lifetime, which allows us to time-gate the noise, and using fast optical switches, we were able to measure a SNR = 7.9(8) (offset-corrected) which is about a factor 2-times higher than what was originally achieved in the hyperfine memory in [56]. The demonstration of non-classical statistics for the read-out remains open. Having identified some of the current limitations, an optimized design of the setup, especially the vapor cell, will help us overcome them in the next iteration of the setup.

By correcting the SNR measured with the attenuated laser for the detection efficiency of a WCP in an HBT setup, we can estimate the expected SNR for other sources. The expected SNR can be estimated by $\mathrm{SNR}_{\mathrm{exp}} = \eta_{\mathrm{source}} \times \mathrm{SNR}/\eta_{\mathrm{HBT}}$, where η_{source} describes the outcoupling efficiency of the source intended to be used. By plugging in the heralding efficiency of the SPDC source, we expect maximally a SNR = 4.7(5) for the same storage time. In order to be able to suppress the uncorrelated noise background, however, the storage time needed to be prolonged from 80 ns to 160 ns. The end-to-end efficiency for the longer storage time is expected to be approximately 71 % of the shorter one, which reduces the expected SNR as well. All together, this results in an upper bound of 3.3(4) for the SNR we could have expected for the HPB memory interfaced with the SPDC source, judging from its performance with the WCPs. These numbers rely on the assumption that the correction of the noise, performed to account for the leakage of the signal EOM, is adequate. Furthermore, there are other factors affecting the SNR, which are not so easy to estimate numerically. The forbidden pump was not used in the measurement with the downconverted photons since it added noise. The total amount of atoms in the $m_I = +\frac{3}{2}$ manifold was therefore reduced, thus decreasing the efficiency of the storage process. Although it is hard to estimate by which factor this effect influences the SNR, it does decrease the above mentioned upper bound. The measurements presented above were also recorded at different atomic temperatures. The higher atomic temperature used for the storage experiment with the downconverted photons can result in higher noise levels and therefore lower the SNR. The considerations above rely on the assumption that the WCPs accurately mimic the heralded single photons. While we modeled the temporal envelope of the pulses to resemble those of the SPDC photons, we have no certainty that the modulator generates transform-limited pulses. In that scenario the WCPs we stored would have a lower bandwidth than the heralded photons, which could result in a reduced performance due to a bandwidth mismatch. For future developments a spectral analysis of the pulses would be helpful.

In the interfaced experiment the read-out is contaminated by noise, so that the

single photon nature of the retrieved light could not be measured. The second-order correlation is not compatible with an incoherent model, suggesting that coherent noise processes are taking place. Due to the unwanted σ^- component present in the control light far off-resonant (about 55 GHz detuned) FWM becomes possible. As we have seen in Ch. 2 FWM can become dominant over EIT even for $\Delta = 0$ if the OD is high enough [164]. However, not only the OD, which in our scenario is moderate, but also the Rb density, which on the other hand is high due to the thin cell, plays a crucial role in determining the dominant process [162]. The high density in our cell contributes to the noise also by limiting the degree of atomic polarization that we can reach, thus leaving atoms in $|s\rangle$, which through stimulated Raman scattering can lead to noise at the signal frequency. Furthermore, the density could also lead radiation trapping to contribute to the detected noise. From the noise characterization measurements we have seen that the detected noise moves with the working point detuning. This behavior is compatible with FWM and stimulated Raman scattering, however not with collisional fluorescence, which remains at the natural transition frequency. The latter can thus be excluded as dominant noise source. All in all, reducing the density while keeping the OD constant would be beneficial for reducing the noise.

Many of the limitations encountered in this proof-of-principle experiment can be improved by designing a vapor cell specifically for this application. From the characterization measurements of the initial state preparation we have learned that a significant amount of atoms remains in the state $|s\rangle$. While the read-out noise caused by imperfect pumping may still allow a decent SNR to be measured when storing WCPs, for the measurements with the downconverted photon as input it represents one of the limitations that prevented us from measuring non-classical photon statistics of the retrieval. As mentioned above, we are most likely limited by radiation trapping in our vapor cell. Due to the short cell length we are forced to work at high Rb densities in order to reach ODs for which the process is efficient. Furthermore, Ar as atomic buffer gas does help in confining the atoms, but lacks the degrees of freedom of a molecular buffer gas, which lead to quenching. To improve the atomic state preparation, a vapor cell filled with a few tens of mbar of N_2 should be used, where the exact buffer gas pressure still needs to be investigated.

The OD is of paramount importance for an efficient storage and retrieval process. The required high temperatures have further drawbacks. More collisions take place, which results in dephasing when occurring during the write/read process [160], and as mentioned above the resulting high number densities are disadvantageous for noise processes. In the next iteration of the experiment longer vapor cells should be considered in order to achieve the same OD at a lower atomic temperature. Considering the spatial magnetic field homogeneity, which I measured at the beginning of my thesis (see Appendix B), we were overly cautious regarding the utilized cell lengths. On the other hand we did not have a vapor cell with dimensions laying between large spectroscopy cells and the micro-fabricated one at our disposal. Assuming the same homogeneity along the optical axis as measured in the vertical direction, simply using a thicker cell of the order of 10 mm in transmission would result in frequency

shifts of at most a few MHz along the dimension of the cell. Alternatively, we could use a miniaturized cell designed to internally route the beam along the vertical axis, while having an optical path of about 10 mm. Various designs have been already realized in micro-fabricated cells where the beam is internally routed by reflections on angled surfaces to increase the path length. These designs were fabricated either by micromachining [268], or directly etching [269] the silicon substrate and then HR coating it, or by locally anodic bonding mirrors into the vapor cell [270].

Another way to improve the end-to-end efficiency is to minimize the losses throughout the setup. While the number of optical elements in the signal arm has been kept as low as possible, most optical components were AR coated, and the mode matching for the fiber couplings optimized, the vapor cell itself remains a major cause of losses. The cell has four uncoated surfaces plus the colored glass filter used to absorb the heating laser, and even though it is specified to have low internal losses, the interfaces cause reflections. Usually, after careful alignment of the angle of incidence, a transmission of scarce 60 % through the cell-filter assembly is reached. By AR coating these surfaces[5] and using index-matching gel between the filter glasses and the vapor cell the transmission could be significantly increased. Using the index gel alone already increases the transmission through the cell assembly effortlessly to 71 % (measured at the D_1 line with unheated atoms). In this way fewer photons are lost on their way in and out of the memory. It would be even more convenient to directly reduce the number of surfaces, for example by using the filter glass itself as material for the windows of the vapor cell. The viability of such an approach, however, should be verified through a dedicated study, in order to assess specific technological challenges such as hermeticity of the cell after bonding and differential thermal expansion, to name a few.

Currently, the memory lifetime is limited due to diffusion of the atoms out of the interaction volume. For broadband storage, large control intensities are necessary to obtain an efficient process. With the available power, beam diameters < 1 mm are a realistic upper bound. In order to prevent the atoms from diffusing out of the volume addressed by the laser, the vapor cell could be designed to match the beam profile. If paired with spin-preserving wall-coatings (some development/research in the fabrication process would be necessary) the atoms would stay in the volume interrogated by the laser. The new limitation of the memory lifetime would then be the lifetime of the state preparation.

One of the next logical steps in the improvement of this memory is to reduce the spurious readout induced by the limited suppression of the control EOM. Inspired by the fast SOA switches we use for the pump lasers, we have been wondering for some time already whether this technique could be transferred directly to a TA – after all the physics behind the two amplifiers is the same. The obvious challenge was to develop a driver capable of switching currents of several A (instead of a few hundreds of mA) within a few ns, which would have meant substantial research and development time

[5] It would be already helpful to coat the easily accessible surfaces. Due to the fabrication process it could be hard or even impossible to coat the interior surfaces of the vapor cell.

for our electronics workshop and thus could not be justified before corresponding proof-of-principle investigation. Very recently, however, a TA at 780 nm in a standard butterfly package (Eagleyard EYP-TPA-0780-03000-4006-BFU09-0000) with pigtailed fiber input for the seed and a collimated free space output became available, along with a matched commercial fast current driver (Aerodiode CCS-HPP). Following this new development, the idea of directly pulsing a TA through its operating current was revived. The maximum CW output of the TA is specified to be 3 W at 4 A. The driver's manufacturer slightly modified the housing for us to prevent clipping the free space output. Furthermore, they adapted the pin-layout to match the one of the amplifier to allow us to directly solder the pins of the package to the driver, thereby keeping the electrical connections at a minimum length, to avoid parasitic capacitance and thus obtain the steepest possible signal edges. The maximum pulsed peak current is 4 A. A typical rise/fall time of $0.5\,\mathrm{ns\,A^{-1}}$ is specified with a minimum pulse duration of 1 ns. The jitter is less than 8 ps. We immediately started testing the paired devices and the first measurements showed highly promising results. When operated in pulsed mode, the TA reaches 6 orders of magnitude of suppression, as can be seen in Fig. 5.16. This is a great improvement compared to the amplitude EOM suppression (cf. Fig. 5.4). The shortest optical pulse that can be generated with approximately the maximum peak power is about 2.5 ns long. The current driver does not allow for pulse shaping. The temporal envelope of the generated control pulses is best described as a rectangle, with the programmed length corresponding to some on/off measure rather than a FWHM. Having demonstrated the feasibility of directly pulsing a TA within the necessary time scale while reaching a far greater suppression, the next step in the development is its integration in the HPB memory setup, which is currently being pursued by GB. By adding up the unwanted read-out during the storage time and the read-out caused by the last reference pulse, the gain in efficiency that would come with a pulsed TA can be estimated. Correcting the measured counts for a storage time of 80 ns for the WCP measurement approximately an additional 1 % could be gained in efficiency.

Currently, all the lasers in the memory setup are free running. For the fast experimental cycle, the intrinsic frequency stability of the laser is sufficient for integrations up to about 10 min. To be on the safe side we monitored the different frequencies and occasionally adjusted them throughout the measurements. Especially for the the second-order correlation measurements, however, where it is important to integrate over longer periods of time for collecting a statistically relevant number of coincidences, the limits of the frequency stability started to become noticeable. Considering more complex setups, where for example multiple memories and single photon sources are interfaced to form quantum network nodes, we would greatly benefit from locked laser frequencies, since the necessary integration time generally increases with the setup's complexity. When working in the HPB regime, locking the laser on an atomic transition becomes more difficult, since the reference cell used for locking needs to experience the same environment as the experimental cell. To that end, we would either need an intermittent locking scheme that is alternated with the storage experiments, which however would reduce the repetition rate, or a

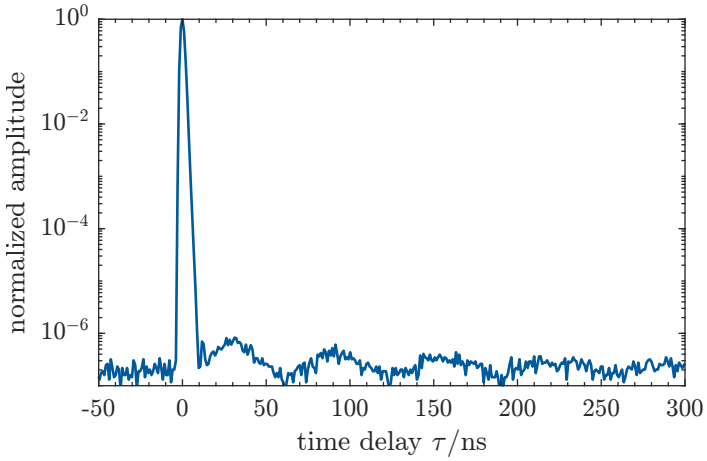

Figure 5.16: Extinction ratio of the pulsed TA after filtering the ASE. The data are binned in 12×81 ps time bins. An extinction of 6 orders of magnitude is reached within a few ns. The extinction ratio is measured with the SNSPD, after the pulses have been appropriately attenuated. We computed the expected dark counts per bin to be $N_{dc} = t_{\text{int}} f_{\text{rep}} f_{dc} t_{\text{bin}} = 1.7$. Comparing this value to the average count floor far away from the ringing of 14.5 counts per bin this measurement results not to be dark count limited.

second vapor cell (or a two chambered cell) would need to be placed in the magnetic field to be used for an atomic lock. With the current electromagnet, however, the latter would be impractical, given the spatial constraints resulting from the magnet's configuration for the most homogeneous field. A commonly used technique consists of locking the lasers to another laser, which is referenced to an atomic transition, through an offset lock by measuring their beat notes. This locking technique was used in the hyperfine memory [56] to keep a constant detuning, corresponding to the ground state splitting, between the signal and control frequencies. Due to the large detunings involved in the HPB regime with respect to the unperturbed lines this, however, might become unfeasible, at least for some transition frequencies. In fact this technique requires fast electronics (with a bandwidth comparable to the frequency splitting), which, depending on the specific transition, could be a limiting factor. A more practical and versatile option, is to lock the lasers on a cavity with an electronic sideband locking scheme, as we did with the source pump laser in Section 3.2.2. For the SPDC source we chose to lock the laser to a passively stable cavity. We used a spacer with an ultra low thermal expansion coefficient to mount the mirrors and isolated the cavity by placing it in vacuum. Depending on the required stability, the demands on the cavity design and price can considerably increase. An alternative would be locking the laser on a transfer cavity – basically the orthogonal approach to a passively stable cavity. Here the cavity is actively stabilized by a piezoelectric actuator. The length of the cavity is locked to a reference laser, which itself is stabilized in frequency, e.g. to an atomic transition. The only requirement to the reference lasers' wavelength is that it be easily separable from the experiment laser, e.g. by a dichroic mirror. Obviously, the cavity needs to be coated to be resonant to both lasers. Once the cavity length is stabilized, the experiment

laser can either be locked directly to the cavity, e.g. with a PDH lock, if the discrete FSR spaced steps give enough frequency resolution, or, alternatively, an electronic sideband lock can be set up to freely choose the frequency to lock onto. Such transfer cavities are used by Tilman Pfau's research group in Stuttgart to stabilize their 1010 nm lasers. In [271] their cavity design is thoroughly described. Since the cavity does not need to be intrinsically stable, the design was chosen to be modular, made mainly out of commercial components, so that the cavity parameters could be easily changed according to the desired requirements. In the beginning of 2022, together with GB, I supervised Demian Ermel who set up such a transfer cavity in the frame of his Bachelor's research project to test its feasibility for future integration in the memory setups.

For future developments of the HPB memory it would be of interest to investigate the usefulness of microwave transitions for the memory scheme. On the one hand, microwave pulses driving transitions between the ground states with $|\Delta m_I| = 1$ could be used to help with the atomic state preparation. The microwave transitions between adjacent ground states of the same m_J-manifold would be interesting to drive for pumping the nuclear spin. Since the respective transitions have similar frequencies, the open question is whether they can be individually addressed. For a description of the microwave transition in a high magnetic field and their transition strength I refer the reader to [168]. On the other hand, nuclear storage could be investigated. Microwave π pulses could be used to transfer the shared coherent excitation to another m_I manifold. Since the nuclear spin relaxes slower than the electronic one, a spin wave between two different m_I states could prolong the memory lifetime.

6 Storing Single Photons in a Zeeman Memory

The second approach we adopt in our laboratory to obtain a clean lambda scheme is to isolate a four-level system by Zeeman-pumping the atomic ensemble. In this approach in which the memory is, in fact, shielded against external fields, we rely on optical pumping and polarization selection rules specific to the ^{87}Rb D_1 line. This scheme allows us to suppress noise and spurious absorption channels, which were limiting factors in the hyperfine memory [56].

While I was working on the SPDC source, Gianni Buser set up the Zeeman quantum memory. Once the memory was operational, GB and I worked together closely to prepare and perform the interfaced measurements with the heralded single-photon source. In this hybrid interconnect we were able to demonstrate, for the first time, storage and retrieval of true single photons in a broadband ground-state vapor-cell quantum memory.

The Zeeman memory is already reported on in full detail in [29] and [254]. In this chapter, in order to give the necessary context for the interfacing experiments that we performed with the SPDC source, I will briefly introduce the working principle and the setup of the Zeeman memory. Subsequently, I will present the storage and retrieval measurements in which the non-classical statistics of the stored photon are preserved in the read-out. Finally, the simulation introduced in Ch. 2 is used to explore the effects of some experimentally controllable parameters. Through comparison of the computed results with the measured data some shortcomings of the simulation are discussed.

6.1 Memory Scheme

The Zeeman memory is operated on the ^{87}Rb D_1 line at 795 nm. The memory scheme is illustrated in Fig. 6.1. Initially, the atomic ensemble is prepared by optically pumping it into the stretched Zeeman ground state $|g\rangle = |F = 2, m_F = 2\rangle$. The storage ground state $|s\rangle = |F = 1, m_F = 0\rangle$ and $|g\rangle$ are coupled to the excited states $|e_1\rangle = |F' = 1, m'_F = 1\rangle$ and $|e_2\rangle = |F' = 2, m'_F = 1\rangle$, isolating two lambda-systems. The incoming, σ^- polarized single photon is coherently mapped to a spin wave between $|g\rangle$ and $|s\rangle$ by a σ^+ polarized control pulse. This memory scheme overcomes FWM noise. Usually, this type of noise is caused by the control pulse coupling off-resonantly to the state $|g\rangle$ [55, 54, 164]. This is avoided by preparing the initial Zeeman state of the atomic vapor in the stretched state, as this state is dark for the

control field. This scheme is specific to the D_1 line of Rb due to the possible values the total angular momentum of the excited states can take. In particular, it does not work on Rb D_2 because of the existence of the $F' = 3$ hyperfine state, which allows off-resonant coupling of the control to $|g\rangle$.

Furthermore, this scheme suppresses parasitic single-photon absorption channels that do not contribute to storage [56]. This occurs when the single photon, instead of being coherently stored, is incoherently absorbed on a transition that would require a selection-rule forbidden mapping by the control to the storage state $|s\rangle$. This can for example happen to atoms initially in the $m_F = 0$ state for hyperfine lambda schemes involving π-polarized light [272]. In the hyperfine memory scheme the state $|F' = 2, m'_F = 0\rangle$ could absorb incoming photons but had no allowed control transitions to store them, lowering the memories efficiency (for details see [254]). These are common problems of memories without control over the Zeeman states and were the limiting factors of the hyperfine pumped memory in [56]. By controlling the Zeeman state and exploiting polarization selection rules both these problem can be overcome [273].

In the isolated four-level system the ground states form a lambda-scheme with each excited state. The pathways to $|e_1\rangle$ and $|e_2\rangle$ interfere destructively [274]. This interference reduces the effective OD of the vapor, but does not lead to absorption without storage. For the operation of the memory the choice of a good working point is essential: It needs to be operated in a regime where one of the two transitions dominates. This sets us in the regime of working point detunings that are lower or comparable to the excited state splitting, but not between the two excited states. A valid detuning rage spans about a GHz red detuned from $|e_1\rangle$ or blue detuned from $|e_2\rangle$.

Working in the high-bandwidth regime is not only appealing for possible future applications, it also is advantageous for reducing the read-out noise. By choosing a single-photon bandwidth much larger than the radiative decay rate of the excited states ($2\pi \times 5.75\,\mathrm{MHz}$ for ^{87}Rb D_1), we are able to discriminate noise by time gating the read-out, significantly suppressing collisional fluorescence noise.

The implemented memory scheme overcomes several major limitations of previous implementations of ground-state memories in atomic vapors, finally allowing us to measure the non-classical signature of the retrieval due to the low contamination through noise counts.

6.2 Experimental Setup

Most devices in the setup of the Zeeman memory are the same as in the HPB memory. I therefore refer the reader to the previous chapter for a detailed description of the setup and components of this experimental implementation. The most significant difference is the available single-photon detector. For the interfacing measurements with the Zeeman memory we still used SPADs to detect single photons.

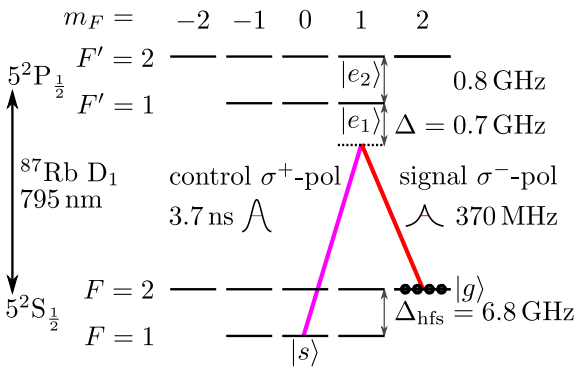

Figure 6.1: Energy level scheme for the Zeeman pumped memory. The atoms are initially prepared in the stretched state $|g\rangle$. Through control over the Zeeman state of the atoms and polarization selection rules a four-level system is isolated. The incoming photon is mapped to a coherent shared excitation between states $|g\rangle$ and $|s\rangle$. By choosing a slight detuning Δ we can work in a regime where the transition to one of the excited states dominates.

Those are the same detectors used for the majority of the measurements presented in Ch. 3 with a jitter of $> 350\,\text{ps}$ and a detection efficiency of about 60(6) %.

As a signal input for the memory, we used downconverted photons from the SPDC source. For the interfaced measurements the source was operated in the double-resonant OPO configuration as described in Ch. 3, including the updates necessary to reduce the crosstalk between the two systems (see specifically Section 3.3). Including the OI in the signal arm and the 60 m fiber delay line, a heralding efficiency of $\eta_\text{h} = 40(4)\,\%$ was measured for a 6.48 ns coincidence window. High quality heralded photons are generated by the source at the desired wavelength and sent to the memory with a $g^{(2)}_\text{input} = 4.21(2) \times 10^{-2}$ and a bandwidth of 370 MHz.

For the Zeeman memory experiment a 75 mm long vapor cell with 19 mm outer diameter is used. The windows are wedged to avoid etalon effects. The cell is filled with enriched ^{87}Rb, with a ^{85}Rb abundance specified as $< 1\,\%$. As buffer gas, 5 Torr (about 6.67 mbar) of N_2 are added to prevent the atoms from moving ballistically and to improve the state preparation through quenching. The walls of the vapor cell are coated in paraffin to preserve the spin polarization of the atoms after wall collisions. The cell is located in a four-layer magnetic shield (Twinleaf MS-1L). A conventional heater (Sacher Lasertechnik) is used to reach an atomic temperature of 50(1) °C. This corresponds to a resonant OD of 25 on the signal transition. The complete setup is sketched in Fig. 6.2.

In the idle state, the Zeeman memory is optically pumped continuously. The preparation of the atoms into the stretched Zeeman state requires two lasers. A circularly polarized pump laser addresses the $|F = 2, m_F\rangle \rightarrow |F' = 2, m'_F = m_F + 1\rangle$ transitions while a repump, operating on the D_2 line, depletes the lower hyperfine state manifold by driving the $|F = 1, m_F\rangle \rightarrow |F', m'_F = m_F + 1\rangle$ transitions. The repump is chosen to be circularly polarized as well to aid the Zeeman pumping and out of experimental convenience. The CW powers at the memory level are 18 mW and 9 mW

for the pump and repump respectively. The two pump laser beams are combined on a dichroic mirror (Semrock RazorEdge LPD02-785RU-25) and are subsequently coupled into the same optical fiber. On the experimental stage, the pump beams have a $1/e^2$ diameter of 1 mm and are overlapped to the control beam under small angle of 2.95(15) mrad with a D-shaped mirror. The initial atomic polarization is estimated to be $> 98\,\%$. Upon the detection of a herald photon the memory sequence is started. At this point the pump lasers are switched off. This task is prioritized by the DDG to allow the atoms in the excited state to relax. Each pump laser is individually switched by a SOA, giving us independent temporal control over the beams. To ensure a good state preparation, a minimum pumping time of 2 µs is enforced for the presented measurement. During this minimum pumping time, detected herald photons are ignored on a hardware level by the DDG. Simultaneously, the DDG also switches off the downconversion source by sending a follower pulse to the Pockels cell causing the rotation of the polarization of the 404 nm light field that pumps the source. The original trigger from the herald SPAD is relayed to the second DDG, which triggers less time-sensitive tasks: passing the TTL pulse on to the time-tagger and initiating the generation of the control pulse. The control pulses are prepared analogously to the HPB memory. The output of a home-built CW ECDL at 795 nm is amplitude modulated by a fiber-integrated EOM. The AWG is programmed in advance with the desired pulse shape. For this experiment the reference pulses are omitted and the AWG pattern consists of a single pulse. Subsequently, the control pulses are amplified with a TA (Toptica BoosTA pro 795 nm). In contrast to the problems encountered in the HPB memory with pulsing the seed of the TA, here this worked smoothly. The ASE filter stage is composed of two 795 nm IFs and two 2 mm etalons (FSR = 51(1) GHz, FWHM bandwidth of 1150(20) MHz). The control pulses are coupled into an end-capped PM fiber and sent to the memory stage. A maximum control peak power of 680(40) mW is estimated at the atoms. For the storage and retrieval measurements the best results were obtained for an estimated power of 190(20) mW.

Headed to the atoms, signal photon and control pulses are outcoupled with orthogonal linear polarizations. They are combined on an AR coated calcite prism after which a quarter-wave plate converts the polarizations into the required counter-rotating circular ones. Signal and control beams are directly focused on the center of the vapor cell with their respective fiber outcouplers. The intention therein is to reduce the number of optical elements in the signal arm. The waists are chosen large enough that the beam sizes do not vary significantly due to divergence within the length of the cell. For the signal (control) a $1/e$ diameter of 480(6) µm (520(6) µm) is measured. A flip mirror before the vapor cell allows measuring the waist of the beams with a beam profiler (DataRay WinCamD-UCD12) at an equivalent distance to the cell's center. After the vapor cell, wave plates linearize and align the polarizations of the two fields to the axis of a second calcite prism that is used to separate the signal from the control. The wave plates before and after the cell are fixed in kinematic rotation mounts for full control over their retardance. This is required to fully exploit the > 80 dB polarization extinction ratio of the prisms.

Spectral filtration is done by three 4 mm (FSR = 25.5(4) GHz, FWHM bandwidth of 550(10) MHz) thick monolithic etalons. They were chosen to match the photon's bandwidth. Each provides a −26 dB suppression at 6.8 GHz from the signal frequency (corresponding to the ground-state splitting in the absence of external magnetic field). In the signal arm, a total CW suppression of −160 dB of the control is achieved, including the polarization filtering. At the same time, a signal transmission of $T = 30(3)\,\%$ through the setup with hot but unpumped atoms is measured. Finally, the memory read-out is detected by two SPADs in HBT configuration. For minimizing counts due to stray light, the detection stage is enclosed in a black foamboard box covered by several sheets of blackout fabric.

6.3 Interfacing Experiments

Figure 6.3(a) shows the arrival-time histogram for a storage and retrieval measurement of SPDC photons. For this measurement, Gaussian control pulses with a FWHM of 4 ns are programmed, resulting in 3.77(4) ns long optical pulses. Their peak Rabi frequency on the $|s\rangle \rightarrow |e_1\rangle$ transition is estimated to be $2\pi\times400(30)$ MHz. The working point detuning of the memory is empirically optimized for SNR and efficiency to $\Delta = -2\pi \times 700$ MHz. The control frequency is set to the target value while measuring it with the wavelength meter and is then adjusted to optimize retrieval. This is done by monitoring the live histogram of the count rates of one of the HBT detectors. The typical rate at which the experiment was repeated is $1.5 \times 10^5\,\mathrm{s}^{-1}$, corresponding to the herald rate of the downconversion source for a pump power of 4.5 mW.

Within an integration time $t_{\mathrm{int}} = 20$ min, a total of $N_{\mathrm{herald}} = 1.60 \times 10^8$ heralds are detected resulting in an equivalent amount of storage attempts. After a predetermined storage time of 160 ns the read-out control pulse is applied. In a 6.48 ns wide read-out window, $N_{\mathrm{ret}} = 4.54 \times 10^5$ counts are accumulated. Two measures to assess the read-out noise were performed. When the signal arm of the source is physically blocked, $N_{\mathrm{noise,mem}} = 2.91 \times 10^4$ counts are recorded in the retrieval time-window. This yields a $\mathrm{SNR}_{\mathrm{mem}} = 14.6(20)$. However, this method does not account for any noise originating from the source or the interface between the two systems. For this reason we perform a second noise estimation where the write pulse is omitted. This scenario is slightly different from the noise measurement performed in the HPB memory, where the read-in pulse was shifted by 20 ns. In that case, the write pulse was delayed by an amount long enough not to store the incoming photon, but still short enough to induce a similar effect on the atoms. By omitting the write pulse we induce a small systematic error as the atomic response during the retrieval time is influenced by the read-in pulse. By omitting the read-in pulse we determine the noise of the whole interfaced system, including the noise originating from crosstalk as well. For this scenario $N_{\mathrm{noise,tot}} = 3.86 \times 10^4$ noise counts were detected within the region of interest. This results in a signal-to-noise ratio of $\mathrm{SNR}_{\mathrm{tot}} = 10.8(15)$. This quantity results in a μ_1 number of $\mu_1 = \eta_{\mathrm{h}}/\mathrm{SNR}_{\mathrm{tot}} = 0.037(6)$. Following Eq. (2.8)

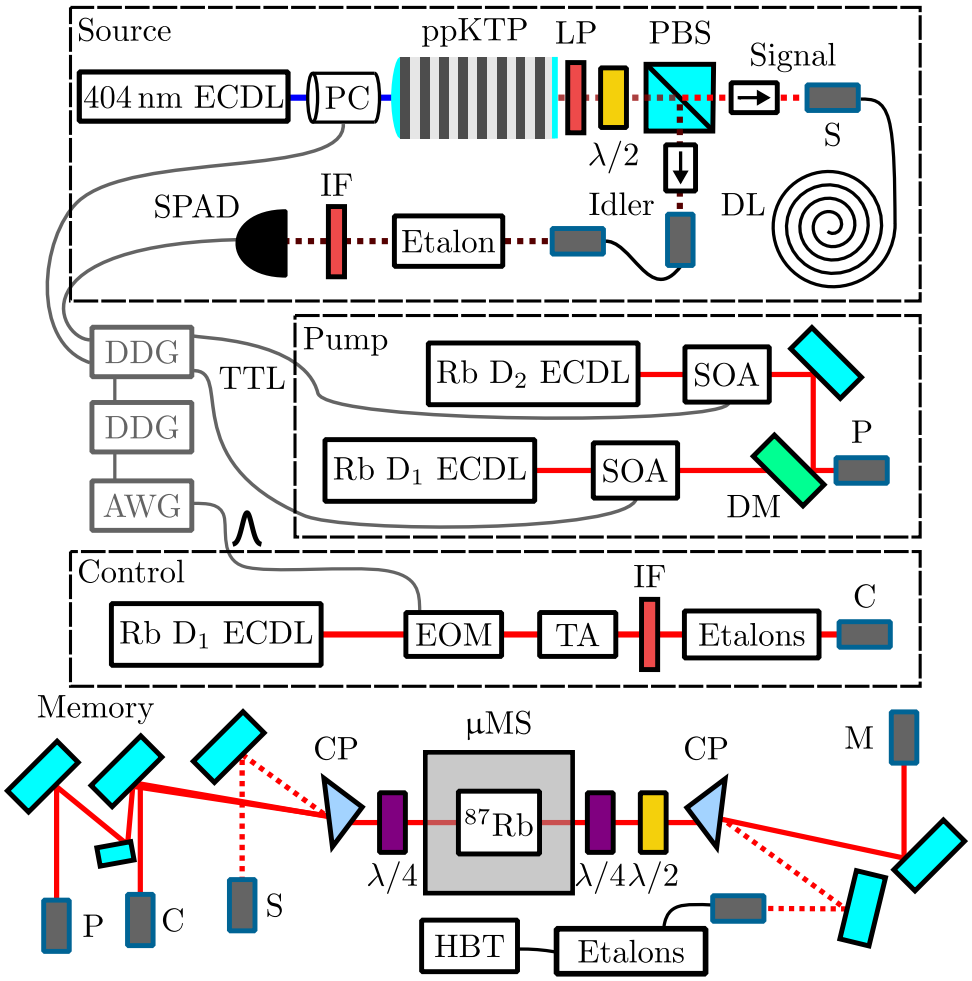

Figure 6.2: Sketch of the setup of the Zeeman memory interfaced with the SPDC source. For a detailed description of the SPDC source see Ch. 3. Used abbreviations: ECDL – external cavity diode laser; PC – Pockels cell; LP – long-pass filter; PBS – polarizing beam splitter; DL – delay line; IF – interference filter; SPAD – single-photon avalanche diode; SOA – semiconductor optical amplifier; DM – dichroic mirror; EOM – electro-optical modulator; TA – tapered amplifier; DDG – digital delay generator; AWG – arbitrary waveform generator; CP – calcite prism; μMS – μ-metal magnetic shield.

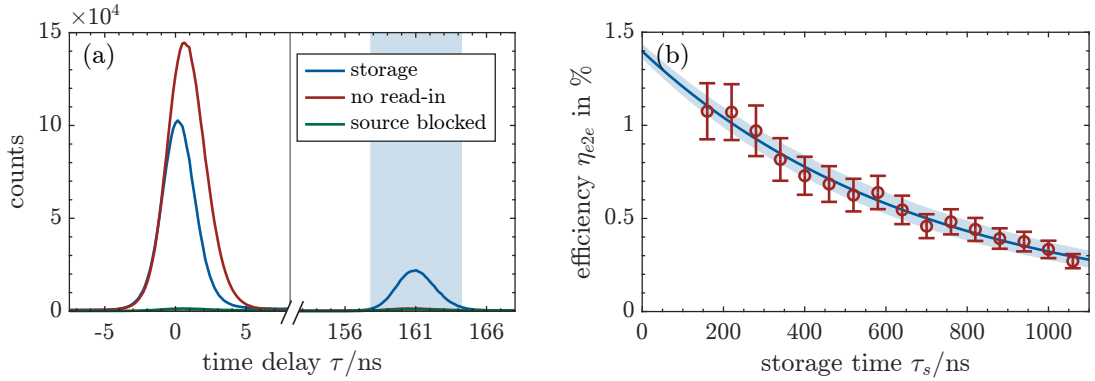

Figure 6.3: (a) Arrival-time histogram of the storage and retrieval measurement of the Zeeman memory combined with the SPDC source. **(b)** Memory lifetime measurement. The drop of efficiency is fit exponentially with $\eta_{e2e} = \eta_{e2e}^{0\,\mathrm{ns}} \exp(-\tau_s/\tau)$. The shaded region corresponds to the 1σ confidence interval of the fit. The presented data was measured a day after the result presented in panel (a) without thoroughly re-optimizing the efficiency. The data were thus scaled by $\sim 15\,\%$ to accurately estimate the internal efficiency of the main result.

the end-to-end efficiency of the Zeeman memory is $\eta_{e2e}^{160\,\mathrm{ns}} = 1.1(2)\,\%$. Computing the conditioned second-order autocorrelation yields a $g^{(2)}_{i:s,s,\mathrm{ret}} = 0.177(23)$. This result constitutes the first successful storage and retrieval of a non-classical state of light from a broadband ground-state memory in atomic vapor. This is the first experimental evidence for the feasibility of warm atomic ground-state memories, putting to rest the prejudice of them being too noisy for preserving the single-photon nature. In our experiment the noise has been successfully reduced to the point that the low $g^{(2)}_{i:s,s}$ at the input remains well below one even after the read-out. The statistical error of the second-order correlation is given by $1/\sqrt{N}$, where N is the number of detected triple coincidences $N = 57$ within the region of interest. The measured triple coincidences are too few to further time-resolve the region of interest with reasonable error bars.

Analogously to the previous chapter, a memory lifetime measurement is performed by repeating the experiment for different storage times (see Fig. 6.3(b)). Each data point represents a storage and retrieval measurement integrated for 5 min. The presented data was collected a day after the main result presented in panel (a) without proceeding to re-optimize the transmission through the setup, which turned out to be slightly lower. The efficiencies were thus scaled, so that the measurement for $\tau_s = 160\,\mathrm{ns}$ matched the value of the main result. An exponential fit with two free parameters yields the $1/e$ lifetime $\tau = 680(50)\,\mathrm{ns}$ and an end-to-end efficiency for zero storage time of $\eta_{e2e}^{0\,\mathrm{ns}} = 1.4(4)\,\%$. Correcting this quantity for the transmission T through the setup yields an internal efficiency of $\eta_{int}^{0\,\mathrm{ns}} = 4.7(14)\,\%$. For longer storage times the $g^{(2)}$ of the retrieval eventually increases. For a storage time of 700 ns the measured value is $g^{(2)}_{i:s,s,\mathrm{ret}} = 0.503(93)$. The lifetime thus corresponds roughly to the storage time for which $g^{(2)}_{i:s,s,\mathrm{ret}} \leq 0.5$.

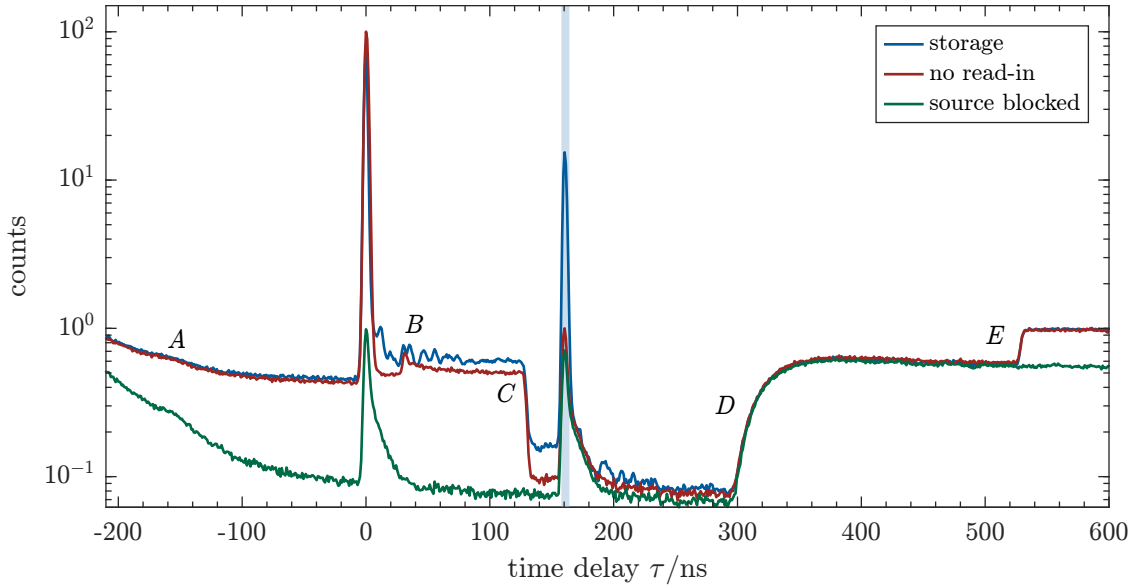

Figure 6.4: Arrival-time histogram of the combined storage and retrieval measurement of the Zeeman memory with the SPDC source. Two different ways to assess the noise are shown. The logarithmic scale allows to recognize small features from the experimental sequence, labeled *A*–*E*, which are described in the main text. The shaded region represents the retrieval time window of 6.48 ns. The data are histogrammed in 1 ns (12 × 81 ps) bins and the y-axis is normalized to the peak of the no read-in curve in the read-out window (8325 counts).

Following the theoretical model in Eq. (2.10), for the measured counts within the retrieval time window, a conditioned autocorrelation of $g^{(2)}_{i:s,s,\mathrm{ret,theo}} = 0.204(29)$ is expected, which is in good agreement with the measured value. For this estimation we used the value $g^{(2)}_{\mathrm{input}}$ obtained from the SPDC source characterization. For the noise, on the other hand, a $g^{(2)}_{\mathrm{noise}} = 2$ was assumed, since the noise counts that accumulated within a reasonable integration time were too few to be statistically evaluated with sufficient significance. Various thermal noise sources are expected. Among those are uncorrelated noise form the heralded source, collisional fluorescence, and control-induced FWM. Leaking control laser light, on the other hand, would result in a $g^{(2)}_{\mathrm{noise}} = 1$. In order to distinguish between these processes, a noise spectrum was recorded by temperature-scanning the last etalon in the filter stage. Collisional fluorescence peaks were indeed present in the spectrum at the expected transition frequencies of the atoms for both $F' \to F = 2$ and $F' \to F = 1$. At the latter frequency the control laser could also be observed. However, those frequency components are well suppressed when the etalon is tuned to the signal frequency. No further peaks were observed in the spectrum, confirming the suppression of FWM. This justifies our assumption of $g^{(2)}_{\mathrm{noise}} = 2$, especially since a coherent noise process would lead, were it present, to a strong underestimation of the estimated $g^{(2)}$ of the retrieval [55]. By plotting the arrival-time histogram of the SPDC photon storage experiment with a logarithmic y-axis, as shown in Fig. 6.4, various small features illustrating the experimental sequence become visible. The three curves represent

the storage and retrieval measurement and the two noise characterizations. Time delay $\tau = 0$ corresponds to the arrival time of the signal photon at the detector for the case in which it is not stored but leaks through the memory. Around label A, when the herald photon has already triggered the memory's DDG, we can see the exponentially decaying fluorescence of the atoms after the pump lasers are switched off. The repump is switched 50 ns after the pump, in order to minimize the population in $F = 1$. After the single photon is stored, unintentional read-out can be observed around label B. It can be recognized as the difference between the storage and the no read-in trace. Unintentional read-out is caused, as in the HPB memory, by leaking control light caused by a combination of the EOM ringing and its limited extinction ratio of about 25 dB. An end-to-end (internal) efficiency of 0.38(5) % (1.26(17) %) can be attributed to the spurious retrieval. The feature at $\tau = 30$ ns that can be seen in both the storage and the no read-in curve just below label B is caused by afterpulsing of the SPADs. The steep edge labeled as C is caused by the Pockels cell switching off the SPDC source (see Section 3.3). Since the Pockels cell is triggered as early as possible, the edge at C consequently represents the minimum storage time for which the uncorrelated noise floor can be suppressed. The shaded region represents the 6.48 ns retrieval time window. The difference of two noise measurement curves within the retrieval window can be attributed to the effect of the read-in pulse on the atoms. At about $\tau = 300$ ns (label D) both pump lasers are switched back on again. The steep edge labeled as E is induced by the Pockels cell rotating the polarization of the 404 nm pump back to generate new photon pairs.

6.4 Simulation

I performed the simulation described in Section 2.3 to compute the internal efficiency as a function of the peak Rabi frequency Ω_{max}. The atomic parameters of the Zeeman memory were used for this computation, as listed in Tab. 6.1. A Gaussian control pulse with a fixed temporal FWHM of 3.77 ns is considered. For each simulation point, the time alignment between signal and control pulse is optimized with a discrete temporal resolution of about 8 % of the control width, corresponding approximately to 300 ps, which is on the order of the detector jitter. To help guide future improvements, experimentally controllable parameters are varied in the simulation.

The simulation was first used to understand the effect of the two-photon detuning Δ_{tp} on the maximum internal efficiency that can be reached for the chosen set of parameters. As can be seen in Fig. 6.5(a), Δ_{tp} has a strong effect on the achievable efficiency. Intuitively, however, one would assume that minimizing the two-photon detuning – thus bringing the memory into resonance – would make the process most efficient. However, following the convention of [131] we define the two-photon detuning as the difference between the control and signal detunings on their respective single-photon transitions, $\Delta_{\text{tp}} = \Delta_c - \Delta_s$. Our experiment is performed in the regime where $\Omega \approx \Delta < \Delta_{\text{hfs}}$. In this regime, a significant, time-dependent light shift

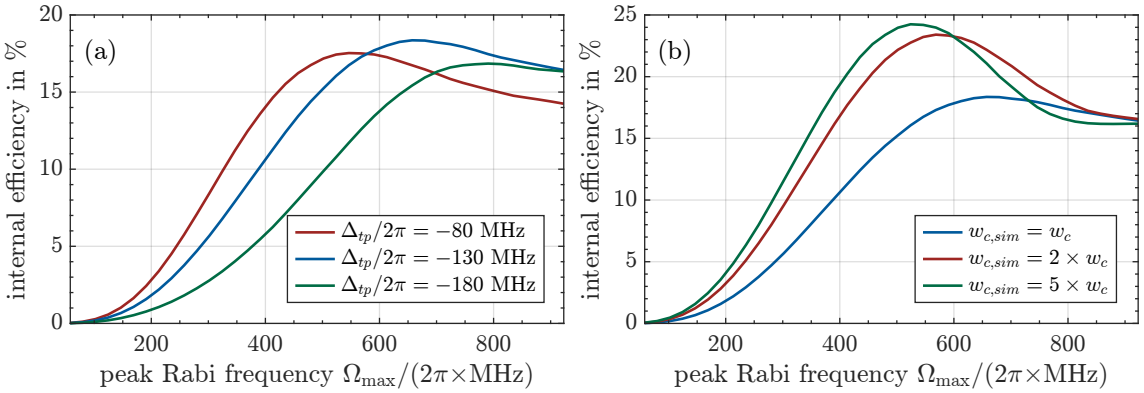

Figure 6.5: Simulation of the internal memory efficiency for the known experimental parameters as a function of the peak Rabi frequency. **(a)** Investigation of the effect of different two-photon detunings on the internal efficiency. **(b)** Dependency of the internal efficiency on the control beam waist. The simulation was run for different control beam waists, where $w_c = 260(3)\,\mu\text{m}$ is the waist of the control in the vapor cell center as used in the experiment. A two-photon detuning of $\Delta_{\text{tp}}/2\pi = -130\,\text{MHz}$ was chosen for these simulations. The blue curve is the same in both panels.

of the energy level is induced by the presence of the control pulse. Following the used definition, however, this appears as a shift of the optimal two-photon detuning from $\Delta_{\text{tp}} = 0$. Modeling the exact experimental scenario is hard since the experimental value of Δ_{tp} is difficult to determine with any certainty. In the experiment the signal frequency is set once, while the control frequency is optimized for a fixed Rabi frequency to yield the maximum end-to-end efficiency. There is no possibility of a direct measurement of the two-photon detuning.

Having implemented a three-dimensional simulation with rotational symmetry around the optical axis allows us to study the impact of different ratios of the beam waists of signal and control beams. Obviously, the control beam needs to be at least as large as the signal, otherwise the area of the signal outside of the control area would just be absorbed by the atoms, instead of being coherently stored. However, it is not straightforward to determine the optimal proportions. For this purpose I ran the simulation for different control beam waists. Figure 6.5(b) shows the simulated performance of the memory for different control waists for the optimal two-photon detuning determined above. The simulation confirms that by choosing a larger control beam waist the efficiency of the process can be increased. With a control beam significantly larger than the signal, a larger fraction of the atoms interacting with the incoming photon experience the optimal control Rabi frequency, leading to a more efficient process overall. This result directly motivated us to choose a control waist of more than double the signal waist for the HPB memory.

The intent of these simulations is to guide future improvements for both memory setups. Taking account of the interdependence of the many involved parameters, we expect that the numerical computations will help us identify the operating regimes with the best performance.

6.5 Discussion

The signal to noise ratio achieved in this implementation demonstrates that ground-state vapor memories operated in a broadband regime are, in fact, not too noisy to preserve the number statistics of a stored photon after retrieval. Even though the Zeeman memory was not optimized for long storage times, its time-bandwidth product of $B = \tau \times \delta\nu = 250(20)$ would be sufficient to directly show enhanced synchronization rates of heralded SPSs.

The efficiency resulting from the measurements shown in Fig. 6.3(a) and Fig. 6.4 falls short of the values predicted by the simulation. This led us to identify some experimental mechanisms that reduce the memory's efficiency not captured by the numerical simulation. The first effect not accounted for is the unintentional read-out, which I mentioned before since it affects both memories. It is caused by the ringing and limited suppression of the EOM inducing the shared excitation to be read out prior to the time window of interest. This effect can be mitigated by using optical switches with a better extinction ratio. To this end, we are investigating the possibility to directly pulse a TA (see discussion of Ch. 5), which so far showed a promising suppression of six orders of magnitude. Another issue not considered in the simulation is that in the experiment the time alignment is limited by the $> 350\,\mathrm{ps}$ jitter of the detector. In [254] GB performed some tests with laser pulses, where the detector jitter can be eliminated, to estimate the effect on the efficiency. Due to the temporal misalignment of signal and control caused by the jitter of the herald SPAD we expect the efficiency to decrease by a factor of $< 0.90(5)$. In the experiment described in this chapter, the time-alignment was limited by the timing resolution of the herald SPAD. By replacing the SPADs with faster detectors, as the SNSPDs we now have at our disposal, which have a timing resolution an order of magnitude better, the jitter would be limited by the electronics. The remaining jitter on the control pulse is expected to be $126(3)\,\mathrm{ps}$. These effects, plus the fact that experimentally determining the two-photon detuning is difficult, make modeling the exact experimental situation difficult. A last technical improvement that was identified to enhance the performance of the current setup is the control laser. A more powerful laser would allow us to widen the control beam waist in order to achieve a more homogeneous interaction region.

Apart from the above mentioned technical improvements, the memory could be made more efficient by increasing the OD. This change is technically straightforward to implement by increasing the cell's temperature. However, this would simultaneously affect other figures of merit, especially the read-out noise. Further, because of the interdependence of the parameters, with a change of the OD the working point detuning and the Rabi frequency would need to be reoptimized. For this task, a prediction of the noise, the fidelity, and the lifetime, which currently are not implemented in the simulation, would be of great help. Furthermore, as we have seen in the HPB memory, increasing the temperature can lead to high atomic densities that become a limiting factor (e.g. for the state preparation).

The lifetime of the Zeeman memory, which was not optimized for storage time, is limited by the atoms diffusing out of the interaction volume. By matching the transverse vapor cell size to the beam profile, the lifetime could be prolonged by keeping the atoms in the region addressed by the control. In order to be useful, the cell walls need to be anti-relaxation coated so that wall collisions do not induce decoherence.

[1] In the simulation d is computed by using the equation for the resonant OD of a two-level system $\mathrm{OD} = \frac{3\lambda^2}{2\pi} n_{Rb} L$. Here, the first term represents the resonant scattering cross section, which depends on the wavelength λ. The atomic number density n_{Rb} is given by $n_{Rb} = \frac{p_{Rb}}{k_B T}$ where p_{Rb} is the rubidium vapor pressure (calculated according to the formula provided by Siddons *et al.* [228]), k_B is the Boltzmann constant and T is the temperature of the atoms. This approach assumes a perfectly pumped scenario.

[2] For the simulation, the peak Rabi frequency $\Omega_{\max}$ is derived from the experimentally available peak power $P_{\max}$ of the control pulse. The expression $\Omega_{\max} = 0.5\sqrt{(\Gamma^2 I_{\max})/(2I_{\mathrm{sat}})}$ is used. The leading factor directly gives the theoreticians Rabi frequency. In this expression, $I_{\max} = \frac{2P_{\max}}{\pi\omega_c^2}$, where the factor of 2 accounts for considering a $1/e^2$ control waist ω_c. I_{sat} represents the saturation intensity.

Symbol	Property	Value
L	Vapor cell length	75 mm
T	Temperature	50 °C
d[1]	Theoretical resonant OD for a two-level system	~ 1200
λ	Wavelength	794.978 nm
Δ_e	Excited state splitting	$2\pi \times 814.5$ MHz
Γ	D_1 decay rate	$2\pi \times 5.75$ MHz
Γ_{hom}	Pressure broadening	82 MHz
Δ_s	Detuning from resonance	$-2\pi \times 700$ MHz
$\Delta_{\text{tp}} = \Delta_c - \Delta_s$	Two-photon detuning	$-2\pi \times 130$ MHz
$\tau_{0,s}$	Time constant of the signal photon envelope	1.2 ns
w_s	Signal waist	240 µm
FWHM_c	Temporal width of the Gaussian control pulse	3.77 ns
$P_{\max}$	Control peak power	190 mW
${\Omega_{\max}}^2$	Theoretical peak Rabi frequency	$\sim \pi \times 415$ MHz
w_c	Control beam waist	260 µm
$\mu_{1,g}$	Relative dipole moment $\|g\rangle \to \|e_1\rangle$	$\sqrt{1/2}$
$\mu_{2,g}$	Relative dipole moment $\|g\rangle \to \|e_2\rangle$	$-\sqrt{1/6}$
$\mu_{1,s}$	Relative dipole moment $\|s\rangle \to \|e_1\rangle$	$-\sqrt{1/12}$
$\mu_{2,s}$	Relative dipole moment $\|s\rangle \to \|e_2\rangle$	$-\sqrt{1/4}$
`nz`	Number of spatial grid points along z	35
`nr`	Number of radial spatial grid points	35
`nt`	Number of temporal grid points	1500
`nv`	Number of velocity classes	21

Table 6.1: Atomic and numerical parameters used for the simulation of the Zeeman pumped memory.

7 Summary and Outlook

The aim of this thesis was to implement, characterize, and interconnect alkali vapor based quantum memories with a compatible SPS in order to demonstrate their viability as key components for quantum networks. The focus was laid on broadband quantum memories implemented in warm Rb vapor and on ways to engineer clean three-level systems. A compatible room-temperature solid-state heralded SPS based on SPDC was developed and improved. Finally, the two systems were successfully interfaced.

Summary We developed an SPDC source optimized for high heralding efficiencies while at the same time being frequency- and bandwidth matched to broadband atomic quantum memories in hot Rb vapor. By adopting a monolithic OPO approach we were able to achieve heralding efficiencies of 40(4) % (53(5) % with a short fiber and no optical isolator) [64], beating the limitations of external cavity designs. In a second iteration of the source design we were able to greatly improve the ease of operability and considerably prolong the life span of the used ppKTP crystal. Quite unexpectedly, we even managed to tune the signal wavelength to the Rb D_2 line. This allowed us to use the same source for testing the performance of the HPB memory with single photons as well.

On the quantum memory side, I successfully demonstrated storage and retrieval of weak coherent pulses at the single-photon level in a novel broadband EIT memory. By using a tesla-order applied magnetic field the atomic energy levels were tailored to a nearly ideal three-level system, reducing possible noise channels. This experiment constitutes, to my best knowledge, the first realization of a quantum memory in a thin MEMS cell [65]. Even though the HPB memory setup can just be seen as a proof-of-principle implementation with room for improvement, a good SNR was achieved that is better than the one measured in the hyperfine memory described in [56]. First interfacing attempts with the heralded SPS were performed. The resulting SNR, however, was not high enough to preserve the non-classical signature of the stored photons. This experiment was helpful in pointing out fundamental improvements that need to be made on the vapor cell to enhance the memory's performance.

Interfacing the heralded SPS with the Zeeman-pumped memory allowed us to demonstrate storage and retrieval of broadband single photons from a ground-state atomic vapor cell memory for the first time. To reach this goal, the memory is prepared in the stretched state by Zeeman optical pumping. By exploiting polarization

selection rules, FWM noise can be eliminated, while operating the memory in a high-bandwidth regime permits us to suppress the tail of the fluorescence noise by time-gating, leveraging the short duration of the signal. The attained $\mathrm{SNR}_{\mathrm{tot}} = 10.8(15)$ resulted in a retrieved $g^{(2)}_{i:s,s} = 0.177(23)$. Judging by the memory's performance, the requirement for an interfacing attempt with a quantum dot SPS can be estimated. The emitters' efficiency needs to correspond to at least $\mu_{1,\mathrm{mem}} = 0.027(4)$ in order to achieve a SNR of 1, assuming it adds no noise.

Technical Improvements The first technical improvements to our experiments are already being implemented at the time of writing. The single-photon detectors limiting the time resolution of the Zeeman memory have been replaced with faster ones for the storage and retrieval measurements of the HPB memory. In order to take full advantage of the SNSPDs' timing, the time-tagger should be updated as well, even though this is not going to influence the jitter of the control pulse. Furthermore, tests where a TA is used as an optical switch to create the control pulses are currently being performed. First results suggest that the achieved suppression should prevent untimely retrieval.

Both memory implementations would greatly benefit from custom-made vapor cells. These cells should be designed to overcome what have proven to be major limitations in the current implementations. For the HPB memory, the state preparation is currently limited by the absence of a quenching buffer gas combined with the high Rb density. Using a thin vapor cell with N_2 as buffer gas would improve the maximum reachable atomic polarization while still confining the atoms to a sufficiently homogeneous region of the magnetic field. Based on the measurements of the homogeneity of the magnetic field, there is good reason to believe that a 10 mm-long cell would be compatible with the current electromagnet. With a longer cell at our disposal, the same OD could be reached at lower Rb number densities. Lower densities further aid the pumping process and additionally reduce noise sources. Both memories would benefit from designer vapor cells where the transverse cell dimension is matched to the beam profile. In this way, diffusion out of the interaction region would not limit the lifetime anymore. In order to exploit the advantages of such a design, anti-relaxation wall coatings are necessary to prevent the depolarization of the alkali spin upon collisions with the cell walls. To fully exploit the scalability potential offered by vapor cell memories, such design changes could be implemented using MEMS technology [52]. Instead of using thin cells in transmission, as we did in the HPB memory, a design where the beam is reflected within the cell would enable longer optical paths through the vapor while keeping the dimensions of the cell small [268, 269, 270]. Furthermore, such cells could be produced using wafer-compatible fabrication techniques, which would allow mass production and reduce costs. MEMS vapor cells have been developed over the last decade and have been successfully used for the realization of compact atomic clocks, magnetometers, and gyroscopes [275, 52].

Synchronization of Single Photons Once a slightly longer memory lifetime (on the order of a few µs) is achieved, it would be interesting to determine the indistinguishability of a photon stored in a vapor memory and one successively emitted from the SPDC source with a HOM interference measurement. This experiment would require the first photon to be stored in the memory until its release upon the detection of the herald photon of the successively generated photon pair. The second photon is then routed through a delay line, instead of going to the memory, so that the coincidences can be measured as a function of the arrival-time delay of the two photons on the beam-splitter. So far, we have always used predetermined storage times in both our memory experiments. By the way they are implemented, however, the read-out could easily be initiated by an external trigger on-demand. The memory would then also read out asynchronously, as is required for most practical applications.

Asynchronous read-out would also allow us to directly synchronize SPDC sources. For the Zeeman memory, even though its storage time currently is not optimized, the time-bandwidth product of $B = 250(20)$, resulting from the high bandwidth, would be sufficient to demonstrate synchronization enhanced coincidence rates.

Quantum Dot Single-Photon Storage A major result and a step forward would be to demonstrate the successful storage and retrieval of a single photon emitted from a semiconductor quantum dot [97, 60]. Such a demonstration would enable many exciting experiments in the field of quantum networking. Considering the progress that has recently been achieved with GaAs quantum dots [98, 61, 62], this goal seems to be within reach. With the current performance of the Zeeman memory and no added noise originating from the source, the expected SNR when interfaced with an arbitrary SPS can be estimated to be $\eta_{\text{source}}/\mu_{1,\text{mem}} = \eta_{\text{source}} \times 37(6)$. With the extraction techniques being explored in InGaAs quantum dots, such as the microcavities that recently led to the excellent end-to-end efficiencies reported in [63], it seems likely that GaAs quantum dots might soon follow with similar efficiencies. Even choosing lower Purcell factors, in order to limit the broadening of the generated photons, such microcavities could enable the interfacing with atomic vapor memories, yielding high SNRs.

Quantum Networking Such an interconnect would open up many possibilities, marking the step from the realization of elementary components to the first simple, yet fundamental, quantum networking applications, which would benefit from deterministic sources due to better scaling of success rates.

A good starting point would be to generate broadband entanglement between two separate memories A and B. By using a 50:50 beam splitter after an SPS, path entanglement between the two spatial modes of the photon is created. The superposition of these two modes can successively be stored by memories in each arm. In this way, the entanglement of the modes is transferred to the two memories. The state of the memories is $|1\rangle_A |0\rangle_B + |0\rangle_A |1\rangle_B$, where 0 and 1 represent the number of

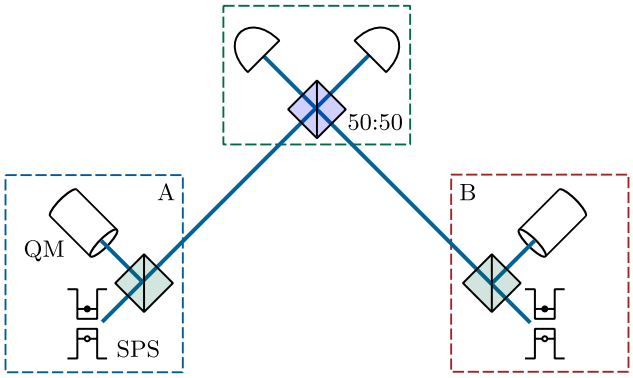

Figure 7.1: Scheme proposed by Sangouard *et al.* in [43] to create entanglement between remote locations. Each node consists of a SPS and a quantum memory (QM).

quanta stored in the memory. This type of entanglement was proven to be as useful as, for example, polarization entanglement [276].

A logical next step would then be to create entanglement between two (remote) memories using a networking scheme, e.g. the Sangouard scheme [43]. The architecture of the Sangouard scheme is visualized in Fig. 7.1. The nodes A and B are each equipped with an SPS that ideally emits a photon when triggered, which is sent to an unbalanced[1] beam splitter. The reflected output mode of the splitter is stored in the quantum memory, while the transmitted mode is sent to a beam splitter in a central station. The detection of a single photon at the central station heralds the storage of the second photon with high probability, due to the unbalanced local beam splitter. Since the measurement erases the which-path information, the two quantum memories are thus entangled, sharing a single excitation. This scheme overcomes the inherent performance limitations of the DLCZ-protocol, since it lacks the fundamental error source of multi-photon generation, which expresses as $|1\rangle_A |1\rangle_B$ error terms corresponding to two full memories and scales proportionally to the entanglement generation probability. This architecture would allow us to exploit the advantages of a hybrid approach.

By adding quantum frequency converters to the outputs of each node leading to the central station in the scheme described above, the distance between the two entangled memories could be increased significantly. Through conversion to the telecom C-band wavelength, the photons could propagate through fibers for significantly longer distances, compared to NIR wavelengths, before absorption losses became an issue. The frequency conversion would furthermore add the benefit that the two memories between which the entanglement is shared, do not necessarily need to work at the same frequency. One could for example operate at 780 nm while the other operates at 795 nm. It has in fact already been demonstrated that QFC can be used to make photons from separate sources more indistinguishable [144].

[1] What transmission to choose for the beam splitter depends upon the exact considered scenario (distance, channel losses, efficiency of the SPS, etc.) and requires optimization for obtaining the best performance of the scheme [43].

The high bandwidth and the room-temperature operation of the developed platform constitute a technologically relevant regime for future experiments. The results reported in this thesis may thus contribute to the development of broadband quantum networks.

Appendix A
Excited State Coefficients

The various constants and g-factors that are needed to compute the Hamiltonian for the ground states or the D excited states can be found in Tab. A.1.

Figure A.1 shows the level structure of the Rb atoms in the HPB regime with the corresponding splittings for an applied magnetic field of about 1.06 T.

A.1 $5\,^2\mathrm{P}_{1/2}$ Term

For the $5\,^2\mathrm{P}_{1/2}$ term the Hamiltonian has the same form as the one of the ground state $5\,^2\mathrm{S}_{1/2}$ (see Eq. (4.8)), merely the magnetic dipole constants are different. The general combination of $|m_J, m_I\rangle$ states corresponds to the one of the ground states, shown in Tab. 4.1. However, the generalized coefficients a_i and b_i reach their final values 1 and 0, respectively, at much lower magnetic fields. The coefficients for the $5\,^2\mathrm{P}_{1/2}$ term are shown as a function of the applied magnetic field in Fig. A.2.

A.2 $5\,^2\mathrm{P}_{3/2}$ Term

For the $5\,^2\mathrm{P}_{3/2}$ term, the quadrupole term in the Hamiltonian is non-zero resulting in the lengthier expression:

$$\hat{H} = \frac{A_{\mathrm{hfs}}}{\hbar^2}\,\hat{\mathbf{I}}\cdot\hat{\mathbf{J}} + \frac{B_{\mathrm{hfs}}}{2I\,(2I-1)\,J\,(2J-1)}\left(3\left(\hat{\mathbf{I}}\cdot\hat{\mathbf{J}}\right)^2 + \frac{3}{2}\left(\hat{\mathbf{I}}\cdot\hat{\mathbf{J}}\right) - I\,(I+1)\,J\,(J+1)\right) + \frac{\mu_B B_z}{\hbar}\left(g_J\hat{J}_z + g_I\hat{I}_z\right)\,. \tag{A.1}$$

The combinations of $|m_J, m_I\rangle$ states can now be composed of up to four states. They can be found in Tab. A.2. Figure A.3 shows how the generalized coefficients vary as a function of the magnetic field. They are arranged as in the low field scenario.

Symbol	Property	Value
$A_{\mathrm{hfs}}^{5\,^2\mathrm{S}_{1/2}}$	Magnetic dipole constant, $5\,^2\mathrm{S}_{1/2}$	$h \cdot 3.417\,341\,305\,452\,145(45)$ GHz
$A_{\mathrm{hfs}}^{5\,^2\mathrm{P}_{1/2}}$	Magnetic dipole constant, $5\,^2\mathrm{P}_{1/2}$	$h \cdot 407.25(63)$ MHz
$A_{\mathrm{hfs}}^{5\,^2\mathrm{P}_{3/2}}$	Magnetic dipole constant, $5\,^2\mathrm{P}_{3/2}$	$h \cdot 84.7185(20)$ MHz
$B_{\mathrm{hfs}}^{5\,^2\mathrm{P}_{3/2}}$	Electric quadrupole constant, $5\,^2\mathrm{P}_{3/2}$	$h \cdot 12.4965(37)$ MHz
g_S	Electron spin g-factor	2.002 319 304 362 2(15)
g_L	Electron orbital g-factor	0.999 993 69
$g_J^{5\,^2\mathrm{S}_{1/2}}$	Fine structure Landé g-factor, $5\,^2\mathrm{S}_{1/2}$	2.002 331 13(20)
$g_J^{5\,^2\mathrm{P}_{1/2}}$	Fine structure Landé g-factor, $5\,^2\mathrm{P}_{1/2}$	0.666
$g_J^{5\,^2\mathrm{P}_{3/2}}$	Fine structure Landé g-factor, $5\,^2\mathrm{P}_{3/2}$	1.3362(13)
g_I	Nuclear g-factor	−0.000 995 141 4(10)

Table A.1: D transition hyperfine structure constants and g-factors for ^{87}Rb as listed in [217].

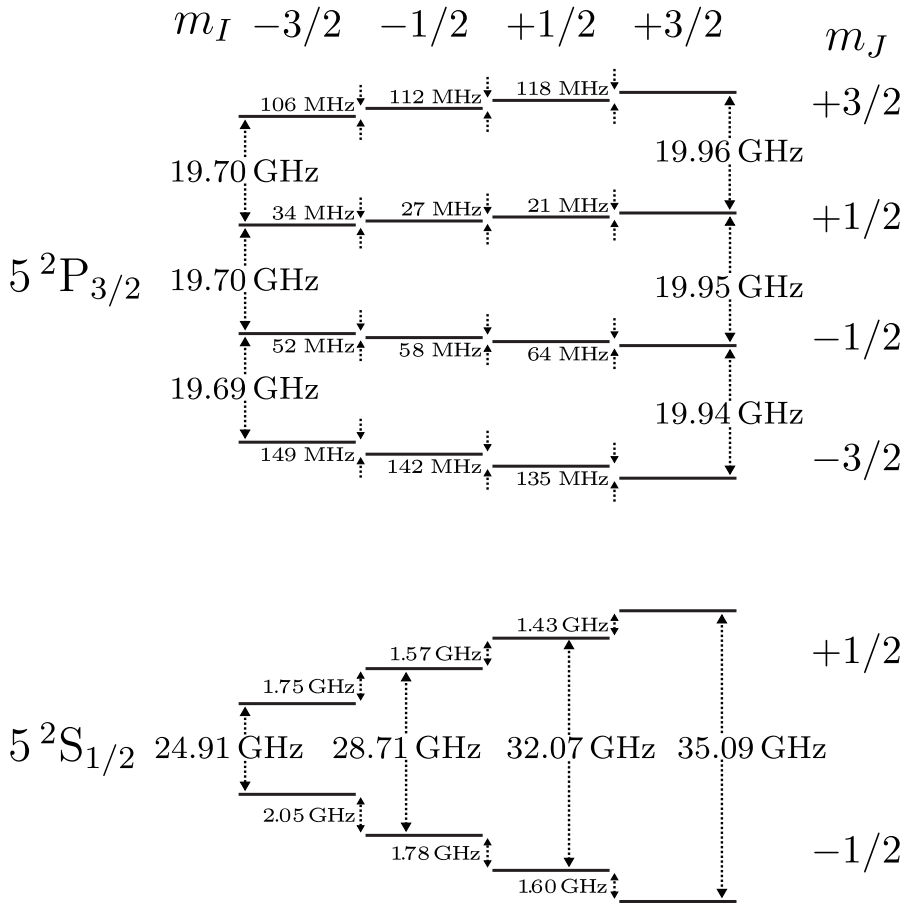

Figure A.1: Energy splittings for the ground and D_2 excited state for an applied magnetic field of 1.06 T.

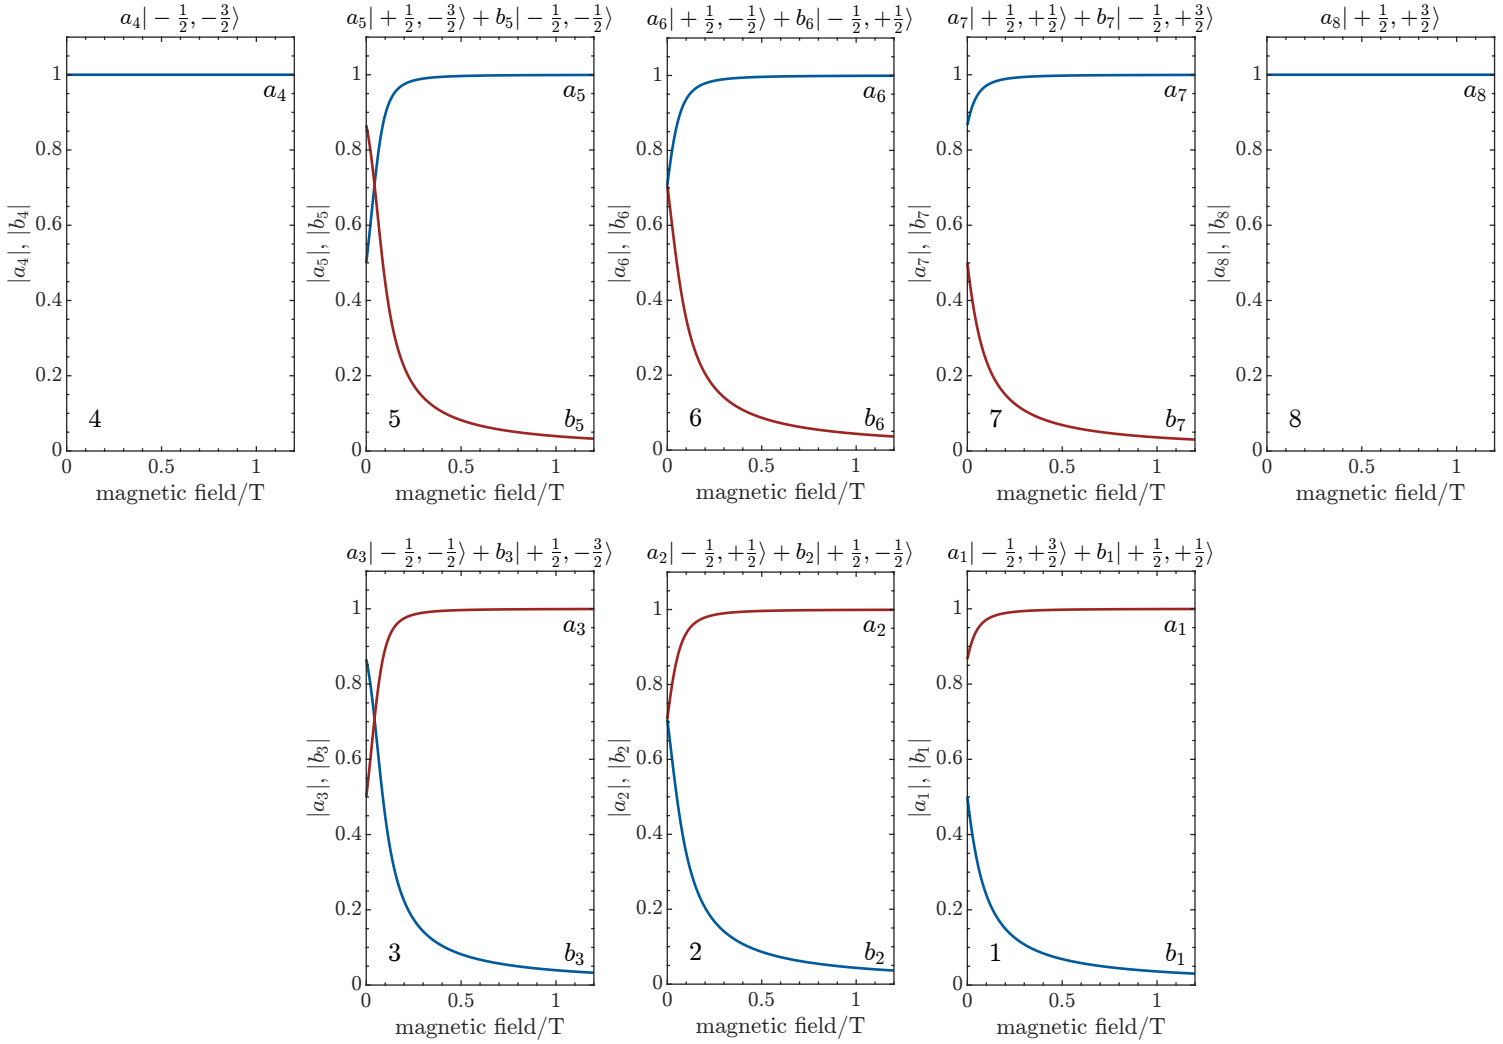

Figure A.2: Generalized coefficients of the $5\,^2\mathrm{P}_{1/2}$ term as a function of the applied magnetic field. The panels are numbered in order of increasing energy of the levels in a high magnetic field.

Level	Weak field $\lvert F, m_F\rangle$	Intermediate regime / General $\lvert m_J, m_I\rangle$	Strong field $\lvert m_J, m_I\rangle$
1	$\lvert 0,0\rangle$	$a_1 \lvert -3/2,3/2\rangle + b_1 \lvert -1/2,1/2\rangle + c_1 \lvert 1/2,-1/2\rangle + d_1 \lvert 3/2,-3/2\rangle$	$\lvert -3/2,3/2\rangle$
2	$\lvert 1,-1\rangle$	$a_2 \lvert -3/2,1/2\rangle + b_2 \lvert -1/2,-1/2\rangle + c_2 \lvert 1/2,-3/2\rangle$	$\lvert -3/2,1/2\rangle$
3	$\lvert 2,-2\rangle$	$a_3 \lvert -3/2,-1/2\rangle + b_3 \lvert -1/2,-3/2\rangle$	$\lvert -3/2,-1/2\rangle$
4	$\lvert 3,-3\rangle$	$\lvert -3/2,-3/2\rangle$	$\lvert -3/2,-3/2\rangle$
5	$\lvert 1,1\rangle$	$a_5 \lvert -1/2,3/2\rangle + b_5 \lvert 1/2,1/2\rangle + c_5 \lvert 3/2,-1/2\rangle$	$\lvert -1/2,3/2\rangle$
6	$\lvert 1,0\rangle$	$a_6 \lvert -1/2,1/2\rangle + b_6 \lvert -3/2,3/2\rangle + c_6 \lvert 3/2,-3/2\rangle + d_6 \lvert 1/2,-1/2\rangle$	$\lvert -1/2,1/2\rangle$
7	$\lvert 2,-1\rangle$	$a_7 \lvert -1/2,-1/2\rangle + b_7 \lvert 1/2,-3/2\rangle + c_7 \lvert -3/2,1/2\rangle$	$\lvert -1/2,-1/2\rangle$
8	$\lvert 3,-2\rangle$	$a_8 \lvert -1/2,-3/2\rangle + b_8 \lvert -3/2,-1/2\rangle$	$\lvert -1/2,-3/2\rangle$
9	$\lvert 3,-1\rangle$	$a_9 \lvert 1/2,-3/2\rangle + b_9 \lvert -1/2,-1/2\rangle + c_9 \lvert -3/2,1/2\rangle$	$\lvert 1/2,-3/2\rangle$
10	$\lvert 2,0\rangle$	$a_{10} \lvert 1/2,-1/2\rangle + b_{10} \lvert 3/2,-3/2\rangle + c_{10} \lvert -3/2,3/2\rangle + d_{10} \lvert -1/2,1/2\rangle$	$\lvert 1/2,-1/2\rangle$
11	$\lvert 2,1\rangle$	$a_{11} \lvert 1/2,1/2\rangle + b_{11} \lvert -1/2,3/2\rangle + c_{11} \lvert 3/2,-1/2\rangle$	$\lvert 1/2,1/2\rangle$
12	$\lvert 2,2\rangle$	$a_{12} \lvert 1/2,3/2\rangle + b_{12} \lvert 3/2,1/2\rangle$	$\lvert 1/2,3/2\rangle$
13	$\lvert 3,0\rangle$	$a_{13} \lvert 3/2,-3/2\rangle + b_{13} \lvert -1/2,1/2\rangle + c_{13} \lvert 1/2,-1/2\rangle + d_{13} \lvert -3/2,3/2\rangle$	$\lvert 3/2,-3/2\rangle$
14	$\lvert 3,1\rangle$	$a_{14} \lvert 3/2,-1/2\rangle + b_{14} \lvert 1/2,1/2\rangle + c_{14} \lvert -1/2,3/2\rangle$	$\lvert 3/2,-1/2\rangle$
15	$\lvert 3,2\rangle$	$a_{15} \lvert 3/2,1/2\rangle + b_{15} \lvert 1/2,3/2\rangle$	$\lvert 3/2,1/2\rangle$
16	$\lvert 3,3\rangle$	$\lvert 3/2,3/2\rangle$	$\lvert 3/2,3/2\rangle$

Table A.2: Table summarizing the general combination of $|m_J, m_I\rangle$ states composing the various energy levels of the ^{87}Rb $5\,^2\mathrm{P}_{3/2}$ term. The levels are listed in order of increasing energy.

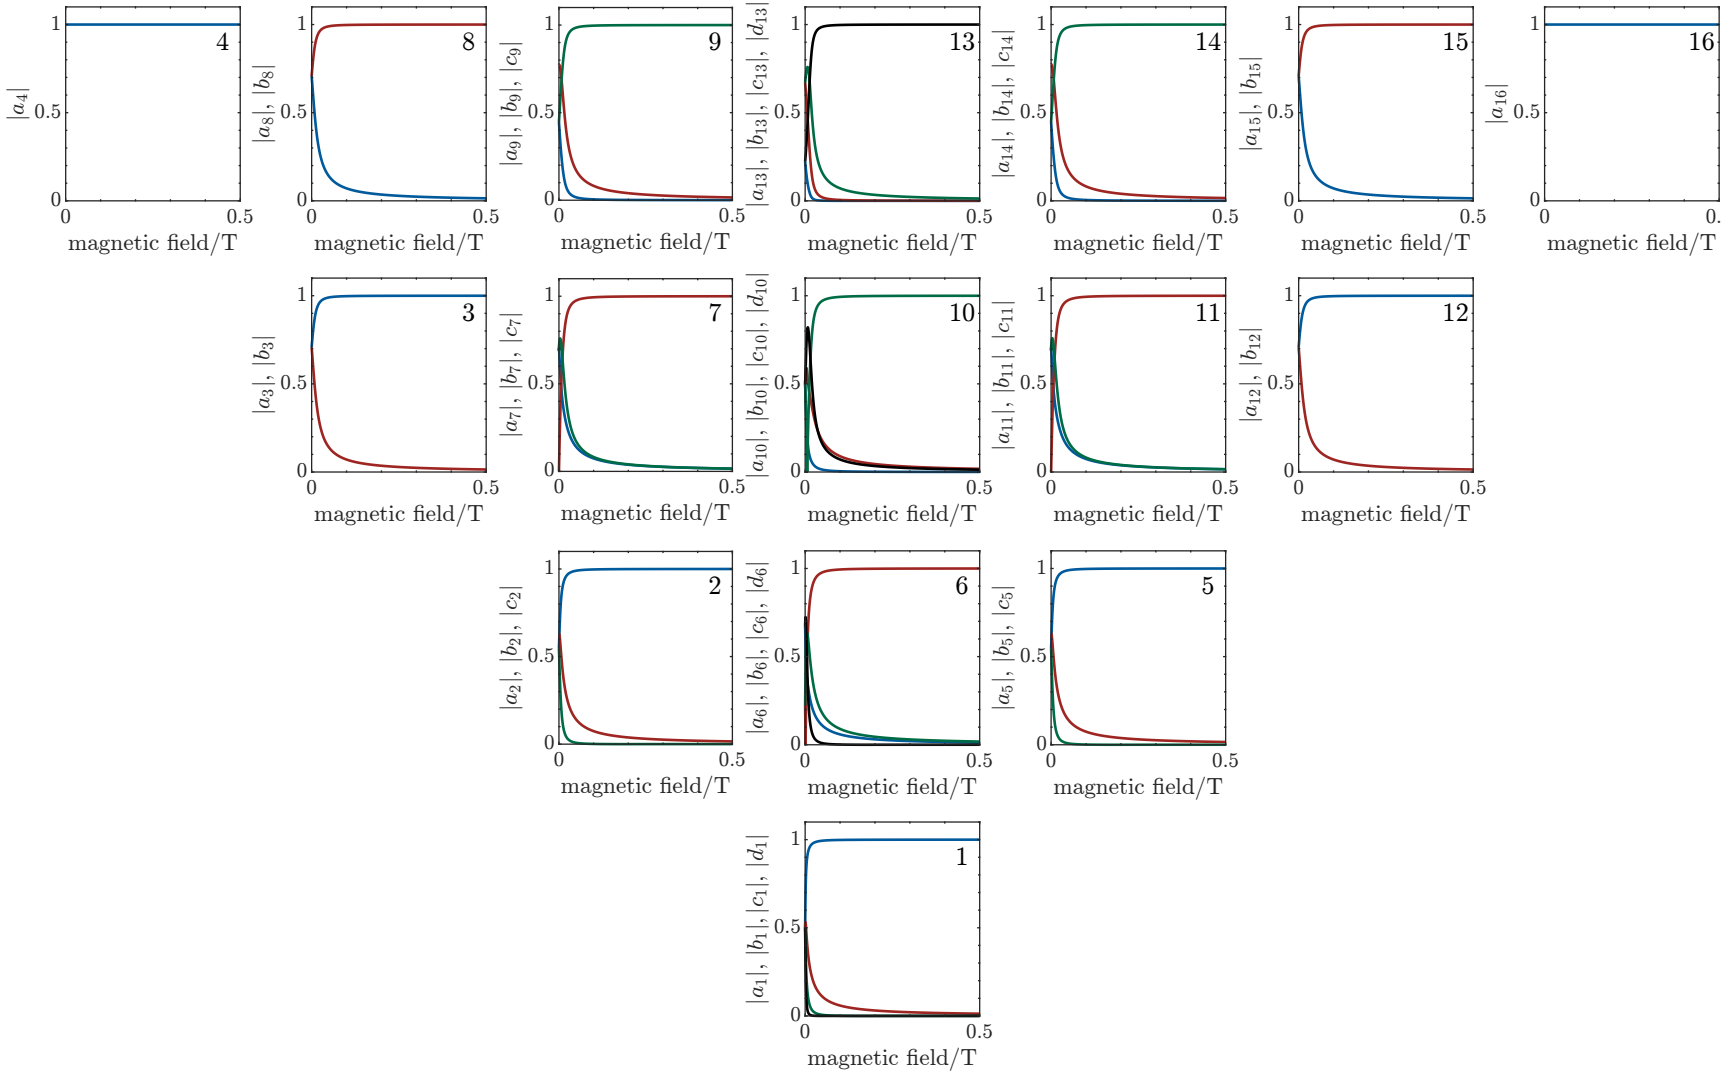

Figure A.3: Generalized coefficients of the $5\,^2P_{3/2}$ term as a function of the applied magnetic field. The panels are numbered in order of increasing energy of the levels in a high magnetic field.

Appendix B
Imaging the Static Magnetic Field

Working in the hyperfine Paschen-Back regime can be advantageous for realizing a quantum memory in hot atomic vapor. However, the applied static magnetic field needs to fulfill specific requirements concerning spatial and temporal homogeneity to ensure that all the atoms within the interaction volume experience the same field throughout the whole storage and retrieval process. This is a prerequisite for the memory to work efficiently. At the beginning of the work leading to my thesis, I characterized the spatial transverse distribution of the static magnetic field generated by the electromagnet by imaging it. For this purpose, I used the imaging techniques described in Ch. 7 and 8 of [168].

Due to the relatively low buffer gas pressure in cell #106 (the cell that has been used for the experiments in Chapters 4 and 5), the lifetime of the state preparation results to be on the order of only 7 µs. The performed imaging characterization is slow. I thus decided to utilize the M2 cell for this characterization, as it has the same geometry and has previously been used by Andrew Horsley for a big portion of his work. The most important properties of this cell are listed in Tab. B.1. Furthermore, for purely technical reasons, the imaging was performed at a field of 0.780(15) T. At this applied field, the ground state splitting of the $m_I = -\frac{3}{2}$ manifold (between states A_4 and A_5) is below 18 GHz, allowing us to use a 6 – 18 GHz microwave amplifier, which has a significantly higher gain than the higher bandwidth amplifier we have at our disposal.

For imaging the static magnetic field, we take a pulsed double resonance (DR) spectrum of the cell. This is done with a pump-probe measurement, similar to a Rabi sequence, with the distinction of a fixed MW pulse duration and a scan of

Property	Value
thickness	2 mm
diameter	5 mm
^{87}Rb abundance	75 %
N_2 pressure	15.3 mbar
Ar pressure	18.7 mbar

Table B.1: Properties of the M2 vapor cell used to characterize the static magnetic field generated with the electromagnet. Values as they are reported in [168].

the MW frequency from shot to shot. Two lasers and a microwave 'antenna' are used for this purpose. The latter is composed of a short-circuited SMA connector (see Section 8.3 in [168]). Both the pump and the probe laser are tuned to be on resonance to the π-polarized transition coupling to the A_5 state.

B.1 Absorption Imaging

For this characterization the technique of absorption imaging is applied (see Ref. [168]). To image the static magnetic field a DR imaging sequence is applied. In a first phase the laser optically pumps the atoms, depleting the ground state it addresses (A_5). Subsequently, a microwave pulse is applied. Its duration is chosen to be approximately that of a π-pulse for the central region of the cell. The sequence concludes with a probe pulse, which is then imaged by the camera.

In order to be able to calibrate the acquired images for spatial variations of probe pulse intensity and stray light, we take reference images and dark images in addition to the actual image. The reference image follows the same optical sequence as the actual image, however, no microwave is injected. It is acquired immediately after the actual image to ensure the same experimental and environmental conditions. Two dark images are captured. These dark images are taken after the actual and reference pulse sequences respectively, with the difference that no probe pulses were applied.

With this method, the difference in optical depth $\Delta\mathrm{OD}_{\mathrm{MW}}$ induced by the microwave pulse can be determined. From the four recorded images the spatial distribution of $\Delta\mathrm{OD}_{\mathrm{MW}}$ can be obtained by applying the following operation to the data

$$\Delta\mathrm{OD}_{\mathrm{MW}} = \ln\left(\frac{I_{\mathrm{act}} - I_{\mathrm{dark1}}}{I_{\mathrm{ref}} - I_{\mathrm{dark2}}}\right) . \tag{B.1}$$

In order to obtain the DR spectrum, the measurement sequence is repeated sweeping the microwave frequency around the hyperfine resonance from shot to shot, while keeping the dark time between the laser pulses and the MW pulse duration fixed. As a first step in the evaluation we determine the resonant microwave frequency as a function of space. Each pixel of the recorded $\Delta\mathrm{OD}_{\mathrm{MW}}$ images is plotted as a function of the frequency sweep, resulting in a DR peak for the evaluated pixel. By fitting this resonance for each individual pixel, we obtain the resonance frequency of the A_4 to A_5 hyperfine transition across the cell. The change in the transition frequency can be linked to local changes of the energy splitting, as the atoms experience slight spatial variations in the static magnetic field. The static magnetic field can be computed numerically from the distribution of the resonant microwave frequency. I do so by comparing the resonance frequency to the eigenvalues of the Hamiltonian used to compute the Breit-Rabi diagrams, described in Ch. 4.

The experiment is controlled by an automated Matlab script slightly adapted from the experiment control scripts AH wrote. Each run of the experimental sequence is

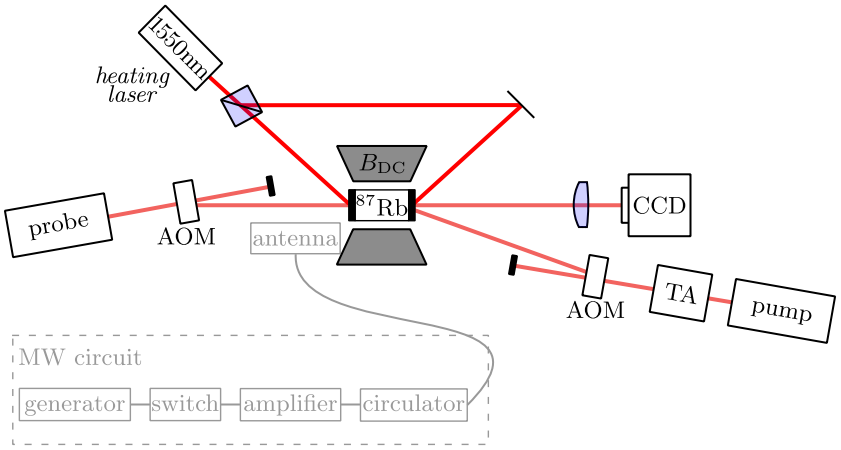

Figure B.1: Imaging setup used to characterize the spatial homogeneity of the static magnetic field. Shown is the optical setup and the microwave circuit.

composed of multiple shots – 500 for the measurement presented further down. From shot to shot, experimental parameters are changed; here the microwave frequency. For this measurement, I averaged the data over 3 identical runs. Usually, the dark images are acquired once a day for each specific setup.

This imaging technique allows us to resolve the transverse direction of the static magnetic field. However, it is important to note that technically, the average of the field over the thickness of the cell along the optical direction is being measured. The position of the cell along the optical axis can be fine-tuned using a linear translation stage by optimizing for the narrowest DR peak. Furthermore, since the experimental sequence runs for several seconds, the obtained field is also an average over time.

B.2 Experimental Setup

For the characterization of the static magnetic field I used a pump-probe setup, shown in Fig. B.1, with two separate lasers, each switched by an AOM (Gooch & Housego 3080-122). Two DFB lasers were used for this purpose. The pump beam is aligned counter-propagating under a small angle to the optical axis. This counter-propagating alignment is chosen to avoid saturating the camera's CCD. In fact, even photons that hit the CCD outside of the exposure time can contribute to noise or even saturation (charge builds up and is read-out at a later point in time). The vapor cell is heated with an IR laser, as described in Ch. 4, with the only difference being that I still used the old heating laser (Seminex 15P-112) here.

Two digital delay generators (SRS DG645 with combinatorial outputs) provide the TTL triggers for all other devices involved in the experimental sequence. The DDGs are interfaced with the computer via TCP/IP. The 'master' delay generator is set to internal single-shot mode, and gets itself a trigger from the Matlab script at the beginning of each shot. This first device relays the triggers to the AOMs of pump

and probe lasers and triggers the 'slave' delay generator. The latter triggers the microwave switch and the CCD camera.

In order to appropriately time the experiment, the various delays need to be measured and considered in the pulse pattern.

Microwave Generation The inset of Fig. B.1 shows the microwave circuit I used for the generation and delivery of the microwave pulses to the vapor cell. For generating the microwave signal I used a synthesized sweeper (Hewlett Packard 8340B, now Keysight Technologies) able to provide an output frequency in the range of 10 MHz to 26.5 GHz, which could be computer-controlled through GPIB. The signal is passed to a switch (American Microwave Corporation SWCH1K-DC40-SK), which is triggered by TTL follower pulses. Finally, the signal is amplified. I used an amplifier (Miteq AMF-2B-06001800-65-35P) which provides 45 dB of gain, at a maximum output power of 35 dBm, in the frequency range of 6 – 18 GHz. The amplifier is protected from backreflections by a circulator (Ditom DMC6018). As an antenna I used a short-circuited SMA jack (Molex 73251-2120). This technique was used before in [168]. Due to its small size, the SMA jack can be placed close to the vapor cell even in the tightly confined geometry of the electromagnet.

The antenna and the IR light might lead to local heating of the cell and thus to the condensation of Rb on the coldest spot of the vapor cell. If Rb droplets build up on the vapor cell windows, it is necessary to clean the cell. In the appendix of [168] an appropriate cleaning technique is described.

Detection For the detection, a scientific monochromatic CCD camera (Allied Vision Guppy Pro F031B) is used. The sensor is composed of 492×656 pixels, each of which has a size of $5.6 \times 5.6\,\mu m^2$. The minimal exposure time corresponds to 75 µs. With a frame rate of 123 fps, the camera is reasonably fast. It, however, constitutes the limiting factor in the repetition rate of the experiment.

B.3 Characterization and Conclusion

Figure B.2 shows an image of the static magnetic field reconstructed from the measured data. For this measurement a microwave center frequency of 17.316 GHz was chosen. A scan of ±12.5 MHz around the center frequency was performed in 50 kHz steps. The initial atomic state was prepared for 50 ms. The duration of the microwave pulse was set to 5 µs. For the probe pulse a duration of 0.5 µs was chosen. The vertical position of the cell was varied until the field decreased both towards the top and the bottom of the cell. Strangely, this was not the case when the cell was placed in the geometric center of the ferromagnetic cores. Surface imperfections and rust building up on the electromagnet may have contributed to this effect.

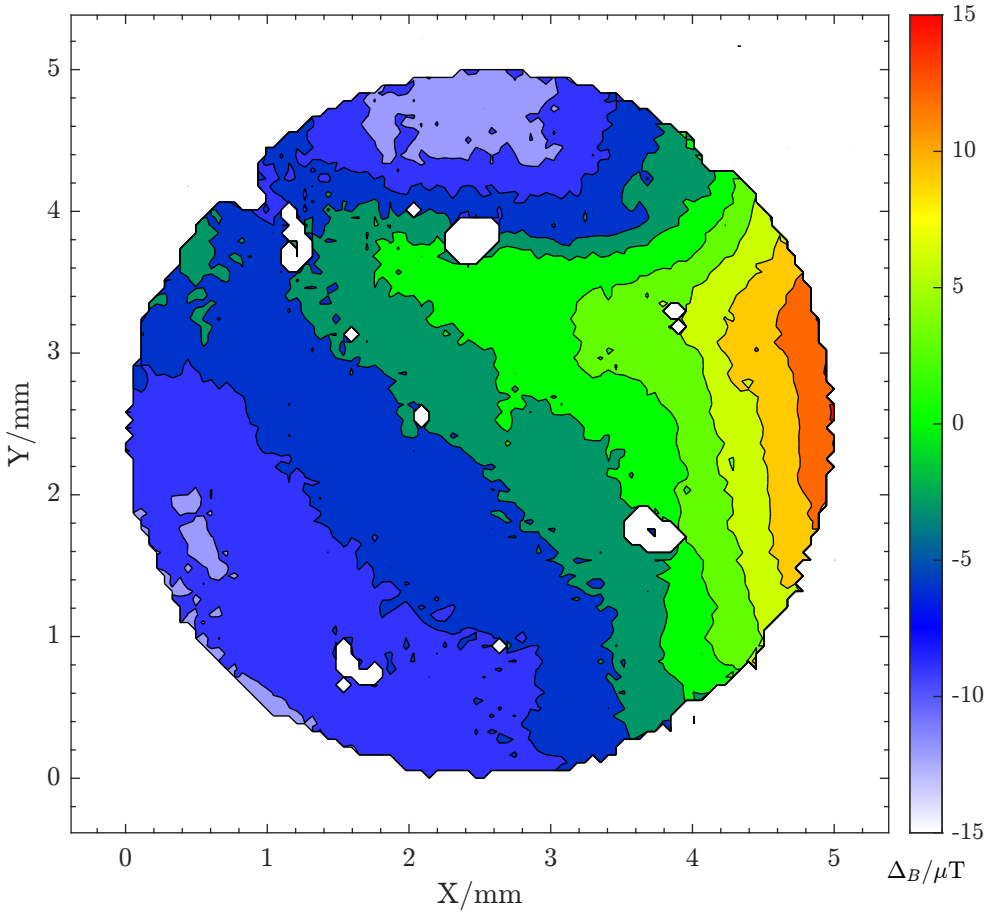

Figure B.2: Spatial distribution of the static magnetic field $B_{DC} = 784.43\,\mathrm{mT} + \Delta_B$ measured with the absorption imaging setup. The contour lines are spaced in 3 µT steps. The whole cell is illuminated by the probe beam. The conversion of the spatial axes from pixels to millimeters has been performed by using the known cell diameter of 5 mm.

The magnetic field varies over 27 µT over the full width of the vapor cell (5 mm). The interaction volume that will be addressed by the signal and control in the storage experiments will be small compared to the whole size of the cell. The beam sizes will be smaller than 0.5 mm, resulting in variations of the magnetic field of less than 5 µT in that specific area. The resulting frequency shifts of the energy levels are well below the frequency fluctuations of the unlocked lasers we use, and should thus not be a concern for the storage experiments.

Appendix C
What's Inside "My" Vapor Cell?

For the quantum memory experiments in the HPB regime we wanted a short, isotopically-enriched cell with a low buffer-gas pressure (for details see Section 4.4.2). The micro-fabricated vapor cells used for previous experiments in the electromagnet all contain a high buffer gas pressure. The Rb atoms are spatially confined by the buffer gas, making these cells well suited for atomic clocks [52] and field imaging applications [235]. However, for the cell with the lowest pressure (M2 cell, see Appendix B) the collisional broadening (also called pressure broadening) due to the buffer gas (513 MHz, see [168]) is already in the order of the Doppler broadening. For our applications this might result in too much broadening. In fact, first attempts performed by Andrew Horsley to observe EIT in that cell remained unsuccessful[1]. We thus opted to start with a lower buffer gas pressure.

Fortunately, our colleagues from Gaetano Mileti's research group at the *Laboratoire Temps - Fréquence* (LTF) at the University of Neuchâtel found some vapor cells they had fabricated in the past, that might have suited our requirements. Among those, two old cells, labeled only as 'BG low press., Rb87', seemed to be promising. Some further characterization of the buffer gas content and ^{87}Rb abundance was necessary. To do so, I visited the LTF on two occasions to perform measurements with the help of Florian Gruet.

C.1 Measuring the Optical Frequency Shift

On my first visit of the LTF in Neuchâtel, back in 2017, we tested 5 old SAM-LAB micro-fabricated vapor cells. The goal was to roughly estimate the buffer gas pressure of these cells by measuring the frequency shift of the optical transitions. For this task we performed a simple linear absorption experiment with a home-built D_2 laser head (v2.2) including an internal Doppler-free reference spectroscopy [277]. For these measurements the analyzed cells were heated to 80 °C in an oven. The frequency shift was measured with respect to the ^{87}Rb $F = 1 \rightarrow F'$ and ^{87}Rb $F = 2 \rightarrow F'$ transitions of the D_2 line. The measured signal was fit with a Voigt profile using the "Multi-peak Fit" analysis tool in the software Igor Pro. Even though documentation about the exact buffer gas composition of the analyzed cells was unavailable, it was still known which buffer gas types were generally used for filling

[1] It has to be noted, however, that these attempts were performed with lower Rabi frequencies compared to the ones now available.

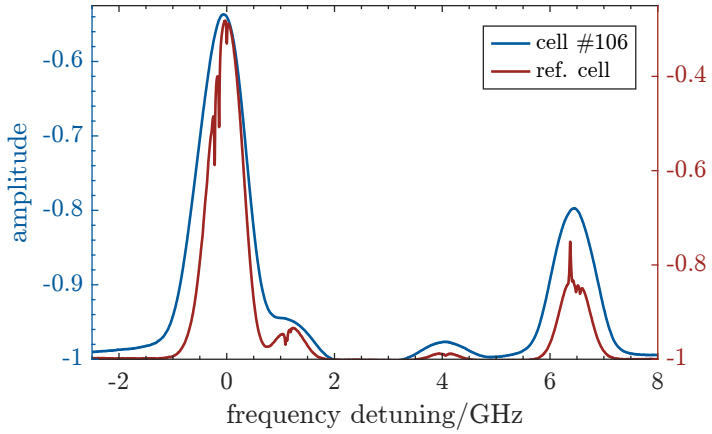

Figure C.1: Relative signals of the linear absorption of vapor cell #106 (blue) and the Doppler-free reference spectroscopy of the laser (red).

Cell	Transition	Frequency shift (MHz)	Estimated buffer gas pressure (mbar)	Avg. pressure (mbar)
#105	^{87}Rb $F = 2 \rightarrow F'$	−66.6	15.3	9.1
	^{87}Rb $F = 1 \rightarrow F'$	−12.7	2.9	
#106	^{87}Rb $F = 2 \rightarrow F'$	−72.5	16.7	9.4
	^{87}Rb $F = 1 \rightarrow F'$	−8.9	2.0	

Table C.1: Buffer gas estimation based on the frequency shift measured with the linear absorption setup.

this type of cells. This knowledge narrowed down the gases in question, leaving us with Ar and N_2 as possible candidates. The shift coefficient for Ar and N_2 is similar and rounded to approximately $-5.8\,\mathrm{MHz\,Torr^{-1}}$ [278].

From the tested cells, #105 and #106 were the only ones that had a buffer gas pressure and an ^{87}Rb abundance compatible with our requirements. Figure C.1 shows the linear absorption spectrum of cell #106 (blue) and the Doppler-free reference spectrum (red). From the fits, the frequency shift relative to the reference cell spectrum could be extracted. The resulting buffer gas pressure for the two promising vapor cells is listed in Tab. C.1.

The results of these measurements are to be treated only as estimates. On the one hand, the fit is not reliable when the absorption peaks are too weak or broadened too much by the buffer gas. On the other hand, the sub-Doppler features of the reference spectrum might also have disturbed the fit. However, these results helped us to pick a cell to be further analyzed in the next step. Since the two cells performed similarly, we decided to proceed with cell #106 for our experiments due to the slightly higher ^{87}Rb enrichment, while its 'sister-cell' (#105) stayed at the LTF in Neuchâtel.

C.2 CPT Hyperfine Spectroscopy

I visited the laboratories in Neuchâtel a second time for a follow-up characterization of the buffer gas content of the vapor cell #106. FG prepared a coherent population trapping (CPT) spectroscopy setup, which we used to investigate the buffer gas content of the cell. This measurement allows us to discriminate the buffer gas species by measuring the buffer-gas-induced frequency shift of the Rb clock-transition frequency, to which Ar and N_2 contribute with opposite sign.

CPT is a phenomenon closely related to EIT. It can be seen as destructive quantum interference between absorption paths. Usually, the phenomenon is referred to as CPT when the strength of the involved fields is comparable. It occurs when the beat frequency between two phase-stable, resonant light fields equals the ground state splitting of the considered atomic species. When the frequency and phase difference of the two light fields equals zero, the atomic population is transferred to a dark-state, a coherent superposition of the ground states. The atoms are 'trapped' because the light fields do not couple to this superposition, which can be detected as an increased transmission through the ensemble. Often the two ground states of the clock transition (the $m_F = 0 \leftrightarrow m_F = 0$ hyperfine transition) are used for CPT measurements due to their first-order insensitivity to magnetic fields. The involved level scheme is depicted in Fig. C.2(a). The unperturbed clock frequency corresponds to the hyperfine splitting $\nu_{\text{clock}} = 6.834\,682\,610\,904\,29(9)\,\text{GHz}$ [279].

The narrow linewidths that can be achieved with CPT make it attractive for metrology applications. The CPT resonance width is inversely proportional to the lifetime of the ground-state coherence. For stable alkali atoms it is determined by external perturbations (e.g. relative phase instability of the light fields, collisions, magnetic field inhomogeneities, and power broadening). Narrow CPT resonances $<$ kHz can be measured even when using lasers with tens of MHz broad linewidths (e.g. DFB), since the relative phase stability, given by the modulation stability (when working with sidebands), is more important than the absolute frequency stability of the fields [280]. A detailed review on CPT can be found in [281].

For this measurement we used a 795 nm laser frequency-locked to a sub-Doppler absorption line of ^{87}Rb. We sent the beam through a phase EOM for generating sidebands at ν_{clock}. This allowed us to simultaneously excite both ground states ($F = 1$ and $F = 2$) of the Rb atoms with two in-phase light fields: with the carrier and one of the sidebands each resonant to one of the transitions. A quarter-wave plate creates the necessary σ polarization[2]. The vapor cell #106 was kept at 80 °C in an oven while a weak magnetic field, for lifting the Zeeman degeneracy, was applied. Meanwhile, a slow sweep of the detuning δ was performed by changing the modulation frequency. The CPT signal is detected on a PD and the error signal generated by a lock-in amplifier allows us to precisely determine the frequency of the CPT resonance. A simple sketch of the setup is shown in Fig. C.2(b). Similar

[2]Note that for π polarized light one of the two transitions would be selection rule forbidden due to: $m_F = 0 \nleftrightarrow m_F = 0$ for $\Delta F = 0$.

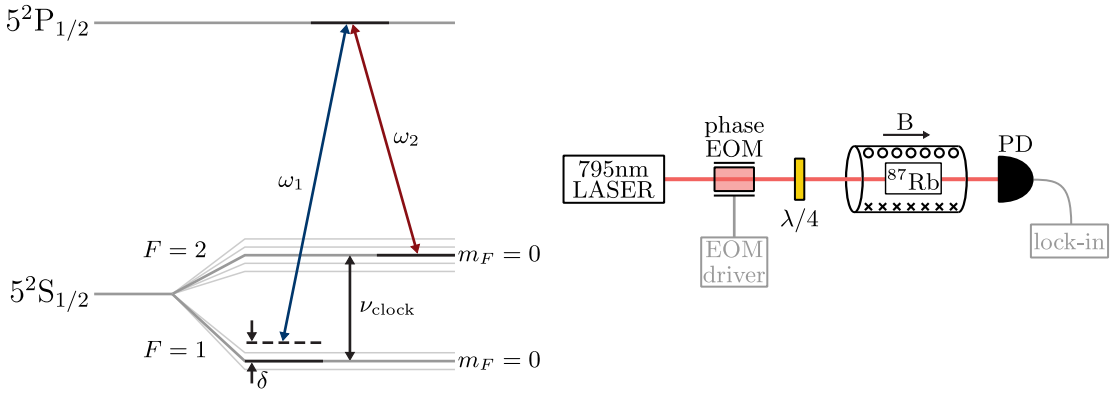

Figure C.2: CPT spectroscopy. **(a)** Simple level scheme for the CPT measurement, as in [275]. **(b)** Sketch of the setup used for the CPT spectroscopy. The light emitted by the laser is sent through a modulator in order to generate sidebands at the clock frequency. The vapor cell is heated and a magnetic field is applied to lift the Zeeman degeneracy. The modulation frequency is slowly swept and the transmission signal is detected with lock-in techniques.

setups for CPT hyperfine spectroscopy are described in detail in Ch. 4 of [282] and in Section 3.18.1.3 of [275].

We measured the clock frequency for different optical input powers to obtain the light-shift graph (see Fig. C.3). Each single data point corresponds to the ν_0 extracted from the CPT error signal. When measuring the CPT resonance frequency, we considered an offset of $-14.25\,\text{mV}$ on the error signal generated by the lock-in detection. Through a linear fit we obtained the extrapolated clock transition frequency at zero optical power, resulting in a shift from the unperturbed Rb hyperfine transition of $-538(8)\,\text{Hz}$. The frequency shift for a given cell temperature T and buffer-gas pressure p is given by the relation $\Delta\nu(T,p) = p\,(\beta + \delta\Delta T + \gamma\Delta T^2)$, where $\Delta T = T - T_0$. The frequency shift coefficients for ^{87}Rb can be found in [283, 284]. For Ar we have $\beta = -59.7\,\text{Hz}\,\text{Torr}^{-1}$, $\delta = -0.32\,\text{Hz}\,\text{K}^{-1}\,\text{Torr}^{-1}$, and $\gamma = -3.5 \times 10^{-4}\,\text{Hz}\,\text{K}^{-2}\,\text{Torr}^{-1}$ specified for a reference temperature $T_0 = 333\,\text{K}$. For a $\Delta T = 21\,^\circ\text{C}$ we obtain a shift coefficient of $-49.93\,\text{Hz}\,\text{mbar}^{-1}$. Assuming that the buffer gas content in the cell was purely Ar, we compute that the buffer gas pressure inside the vapor cell #106 is 10.78(16) mbar. This result is compatible with the estimate resulting from the shift of the optical transitions measured in Section C.1.

The buffer gas could also be composed of a mixture of atomic species. N_2, however, the other plausible candidate that has been used for this type of cells, has a (temperature independent) shift coefficient $\beta = 546.9\,\text{Hz}\,\text{Torr}^{-1} = 410.21\,\text{Hz}\,\text{mbar}^{-1}$ [283, 284]. Being about an order of magnitude larger and with opposite sign with respect to the coefficient of Ar, any (significant) combination of Ar and N_2 would result in a much higher buffer gas pressure, becoming incompatible with our previous measurement from Section C.1. We thus conclude that the cell must contain only Ar as buffer gas.

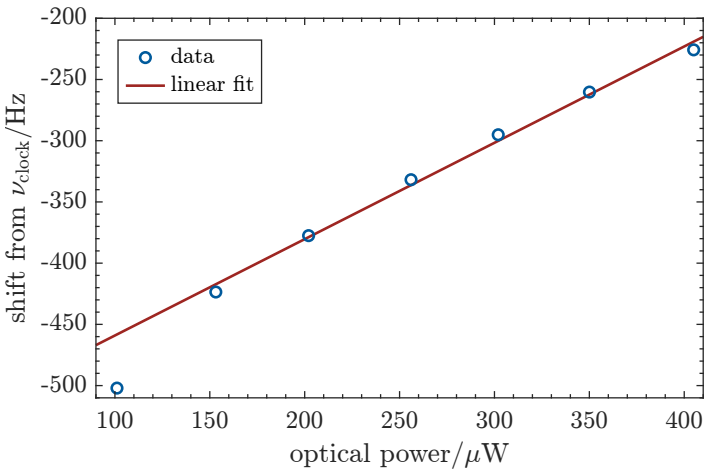

Figure C.3: Light-shift graph for vapor cell #106. The frequency shift of the clock transition is plotted as function of the laser power. The shift at zero power is extrapolated with a linear fit. The first data point has been excluded from the fit due to a weak and noisy CPT signal. The fit yields a shift of $-538(8)$ Hz for zero power.

Appendix D
Absolute Frequency Reference – Rb D_2 Line

The atomic hyperfine structure constitutes an ideal absolute frequency reference, useful, for example, to lock the frequency of a laser or, as in our case, to convert the x-axis of recorded spectra. In order to do so, it is necessary to be able to resolve the hyperfine splittings of an atomic line spectrum. One technique is the so called saturation spectroscopy described in Chapter 8.3 of [251]. In a nutshell, a strong pump beam and a weaker probe propagate in opposite directions through the vapor cell. The pump saturates the absorption of the atoms. Only the atoms with zero velocity component along the optical axis 'see' both beams at the same frequency, due to the Doppler effect. This effect reduces the absorption for this class of atoms resulting in a narrow peak in the transmission of the probe. The same effect happens at the frequency halfway between each transition pair, these features are dubbed 'crossovers'. The resulting spectrum is composed of Lamb-dips within a Doppler broadened spectrum.

Figure D.1 shows a typical example of Doppler-free saturation spectroscopy on the D_2 line of natural Rb recorded with the signal laser. To cleanly resolve the sub-Doppler features the laser frequency was slowly scanned (approximately 0.5 Hz) and the oscilloscope was set up to record in averaging mode. With our spectroscopy setup, we are able to resolve and identify 13 pronounced peaks. In Tab. D.1 the transitions and crossovers with their respective frequencies are listed in the order of increasing scan time. The transition frequencies are expressed as detuning from the line center as defined in ElecSus (see [229, 231]). A (slightly) quadratic fit is performed to linearize the frequency axis. Once the axis is converted from a time to a frequency axis, it can be applied to other simultaneously recorded spectra. In order to have an absolute frequency reference most of the lasers in our laboratory are equipped with such a reference spectroscopy.

An analogous plot of the Doppler-free D_1 line and the transition and crossover frequencies can be found in the appendix of GB's thesis [254].

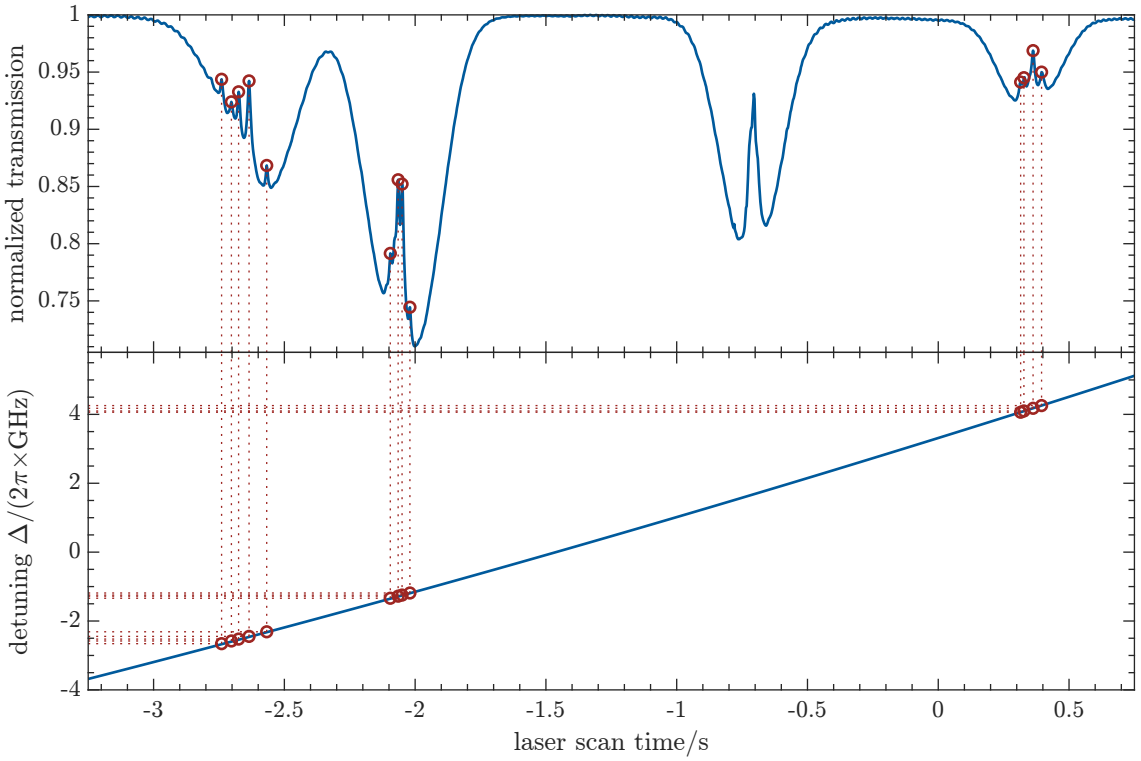

Figure D.1: (a) Doppler-free saturation spectrum of the D_2 line of natural Rb. The distinct/prominent sub-Doppler features used for the conversion of the time axis are highlighted by a circle. **(b)** (Slightly) quadratic fit to convert the time axis recorded with the oscilloscope to a frequency axis that can be used for other spectra measured simultaneously.

Transitions	$\Delta/(2\pi \times \text{GHz})$
^{87}Rb $F = 2 \rightarrow$ CO 12	-2.658
^{87}Rb $F = 2 \rightarrow F' = 1$	-2.580
^{87}Rb $F = 2 \rightarrow$ CO 13	-2.525
^{87}Rb $F = 2 \rightarrow$ CO 23	-2.446
^{87}Rb $F = 2 \rightarrow F' = 3$	-2.313
^{85}Rb $F = 3 \rightarrow$ CO 23	-1.339
^{85}Rb $F = 3 \rightarrow$ CO 24	-1.278
^{85}Rb $F = 3 \rightarrow$ CO 34	-1.247
^{85}Rb $F = 3 \rightarrow F' = 4$	-1.186
^{87}Rb $F = 1 \rightarrow$ CO 01	4.062
^{87}Rb $F = 1 \rightarrow F' = 1$	4.098
^{87}Rb $F = 1 \rightarrow$ CO 12	4.177
^{87}Rb $F = 1 \rightarrow F' = 2$	4.255

Table D.1: List of the distinct transitions and crossover (CO) peaks and their frequency detuning with respect to the D_2 line center as necessary for the fits in ElecSus.

Appendix E
Autler-Townes Splitting

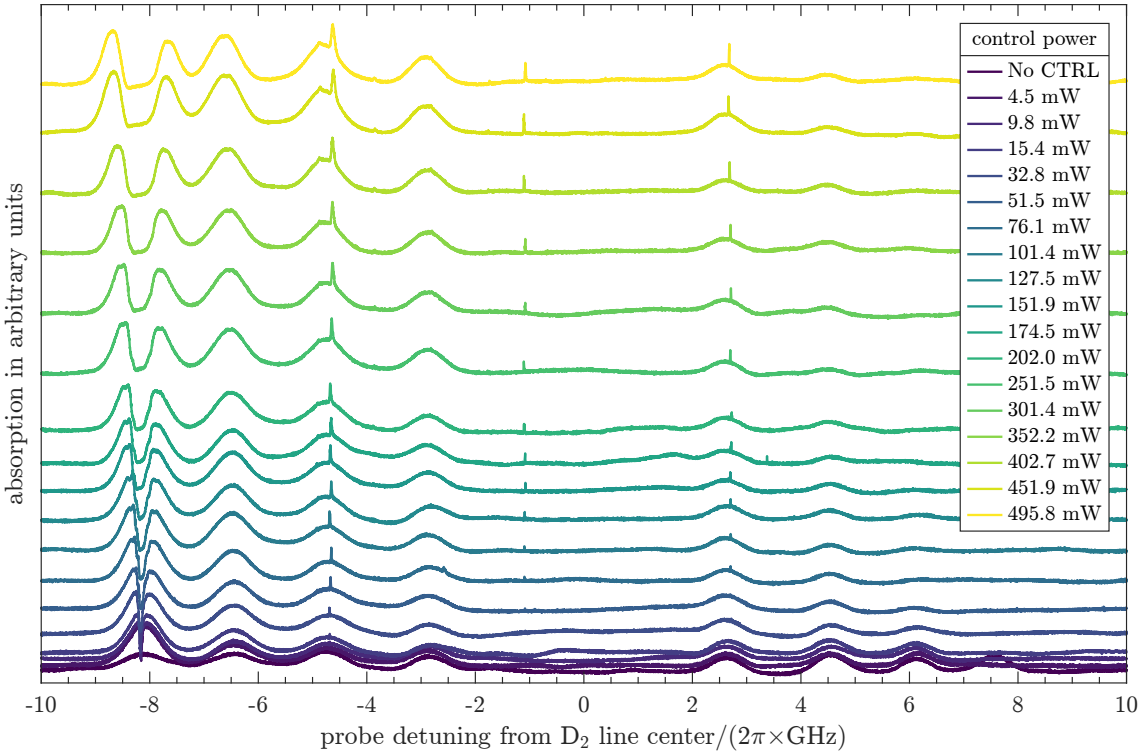

Figure E.1: Experimentally measured absorption of the probe beam as a function of the detuning for different control powers. The transition from EIT to the AT regime can be observed. The different traces are offset along the vertical axis by a value proportional to the their control power. This figure is an alternative representation of the data shown in Fig. 4.11.

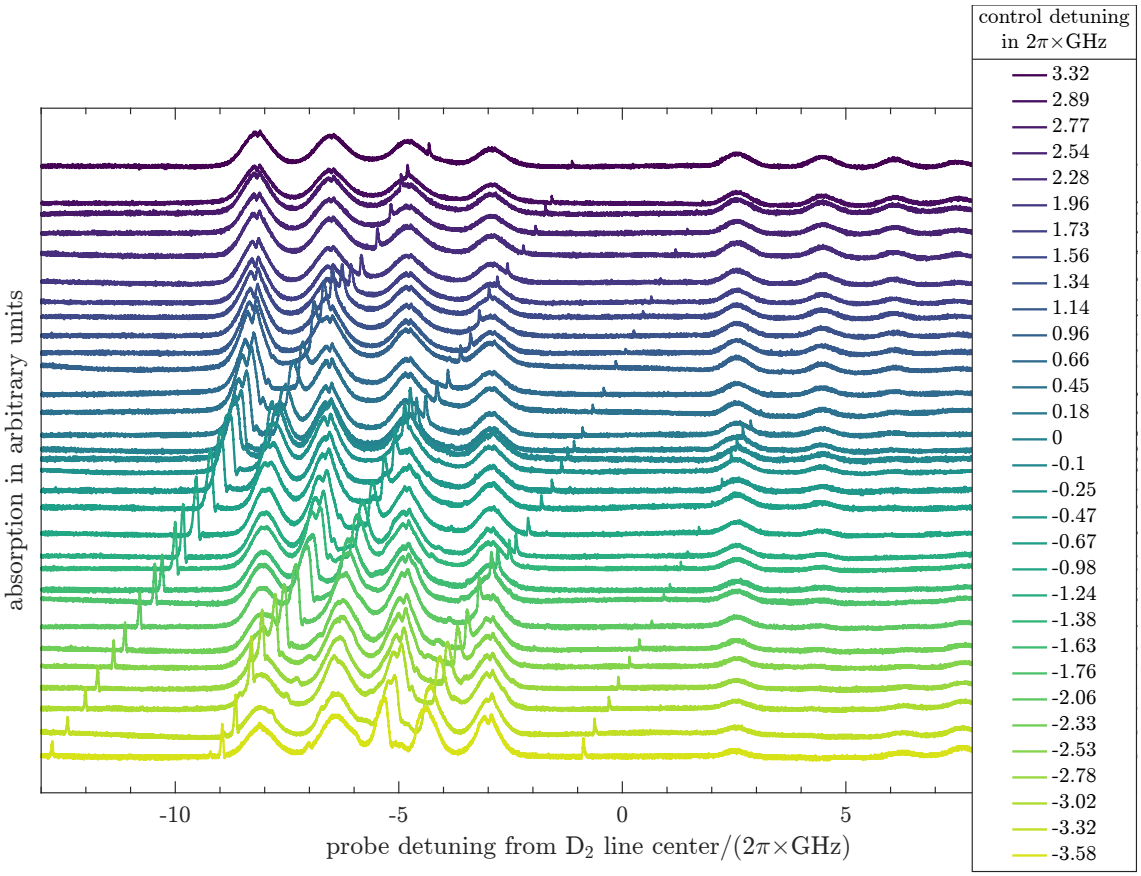

Figure E.2: Avoided crossings in the HPB regime can be recognized in the plot of the absorption as a function of the probe and control detunings. A vertical offset is added to the traces in proportion to the control field detuning. The control detuning increases towards the top of the image. The legend shows the control detuning in GHz. This figure is an alternative representation of the data shown in Fig. 4.12.

List of Figures

List of Tables

Bibliography

[1] S.-H. Wei, B. Jing, X.-Y. Zhang, J.-Y. Liao, C.-Z. Yuan, B.-Y. Fan, C. Lyu, D.-L. Zhou, Y. Wang, G.-W. Deng, H.-Z. Song, D. Oblak, G.-C. Guo, and Q. Zhou, *Towards real-world quantum networks: A review*, Laser & Photonics Reviews **16**, 2100219 (2022).

[2] A. I. Lvovsky, B. C. Sanders, and W. Tittel, *Optical quantum memory*, Nature Photonics **3**, 706 (2009).

[3] C. Simon, M. Afzelius, J. Appel, A. B. de la Giroday, S. J. Dewhurst, N. Gisin, C. Y. Hu, F. Jelezko, S. Kröll, J. H. Müller, J. Nunn, E. S. Polzik, J. G. Rarity, H. D. Riedmatten, W. Rosenfeld, A. J. Shields, N. Sköld, R. M. Stevenson, R. Thew, I. A. Walmsley, M. C. Weber, H. Weinfurter, J. Wrachtrup, and R. J. Young, *Quantum memories*, The European Physical Journal D **58**, 1 (2010).

[4] F. Bussières, N. Sangouard, M. Afzelius, H. de Riedmatten, C. Simon, and W. Tittel, *Prospective applications of optical quantum memories*, Journal of Modern Optics **60**, 1519 (2013).

[5] M. Afzelius, N. Gisin, and H. de Riedmatten, *Quantum memory for photons*, Physics Today **68**, 42 (2015).

[6] K. Heshami, D. G. England, P. C. Humphreys, P. J. Bustard, V. M. Acosta, J. Nunn, and B. J. Sussman, *Quantum memories: Emerging applications and recent advances*, Journal of Modern Optics **63**, 2005 (2016).

[7] M. D. Eisaman, J. Fan, A. Migdall, and S. V. Polyakov, *Invited review article: Single-photon sources and detectors*, Review of Scientific Instruments **82**, 071101 (2011).

[8] H. J. Kimble, *The quantum internet*, Nature **453**, 1023 (2008).

[9] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead*, Science **362**, eaam9288 (2018).

[10] N. Gisin and R. Thew, *Quantum communication*, Nature Photonics **1**, 165 (2007).

[11] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, *Quantum cryptography*, Reviews of Modern Physics **74**, 145 (2002).

[12] F. Xu, X. Ma, Q. Zhang, H.-K. Lo, and J.-W. Pan, *Secure quantum key distribution with realistic devices*, Reviews of Modern Physics **92**, 025002 (2020).

[13] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge University Press) (2010).

[14] R. Van Meter and S. J. Devitt, *The path to scalable distributed quantum computing*, Computer **49**, 31 (2016).

[15] T. Häner, D. S. Steiger, T. Hoefler, and M. Troyer, *Distributed quantum computing with QMPI*, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '21, pp. 1–13 (Association for Computing Machinery, New York, NY, USA) (2021).

[16] L. K. Grover, *Quantum telecomputation*, arXiv:quant-ph/9704012 (1997).

[17] D. Cuomo, M. Caleffi, and A. S. Cacciapuoti, *Towards a distributed quantum computing ecosystem*, IET Quantum Communication **1**, 3 (2020).

[18] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, and P. Walther, *Demonstration of blind quantum computing*, Science **335**, 303 (2012).

[19] V. Giovannetti, S. Lloyd, and L. Maccone, *Quantum-enhanced measurements: Beating the standard quantum limit*, Science **306**, 1330 (2004).

[20] Z. Zhang and Q. Zhuang, *Distributed quantum sensing*, Quantum Science and Technology **6**, 043001 (2021).

[21] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, *A quantum network of clocks*, Nature Physics **10**, 582 (2014).

[22] D. Gottesman, T. Jennewein, and S. Croke, *Longer-baseline telescopes using quantum repeaters*, Physical Review Letters **109**, 070503 (2012).

[23] B. Jing, X.-J. Wang, Y. Yu, P.-F. Sun, Y. Jiang, S.-J. Yang, W.-H. Jiang, X.-Y. Luo, J. Zhang, X. Jiang, X.-H. Bao, and J.-W. Pan, *Entanglement of three quantum memories via interference of three single photons*, Nature Photonics **13**, 210 (2019).

[24] M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, S. Wehner, and R. Hanson, *Realization of a multinode quantum network of remote solid-state qubits*, Science **372**, 259 (2021).

[25] S. L. N. Hermans, M. Pompili, H. K. C. Beukers, S. Baier, J. Borregaard, and R. Hanson, *Qubit teleportation between non-neighbouring nodes in a quantum network*, Nature **605**, 663 (2022).

[26] M. L. G. Puigibert, M. F. Askarani, J. H. Davidson, V. B. Verma, M. D. Shaw, S. W. Nam, T. Lutz, G. C. Amaral, D. Oblak, and W. Tittel, *Entanglement and nonlocality between disparate solid-state quantum memories mediated by photons*, Physical Review Research **2**, 013039 (2020).

[27] N. Maring, P. Farrera, K. Kutluer, M. Mazzera, G. Heinze, and H. de Riedmatten, *Photonic quantum state transfer between a cold atomic gas and a crystal*, Nature **551**, 485 (2017).

[28] K. T. Kaczmarek, P. M. Ledingham, B. Brecht, S. E. Thomas, G. S. Thekkadath, O. Lazo-Arjona, J. H. D. Munns, E. Poem, A. Feizpour, D. J. Saunders, J. Nunn, and I. A. Walmsley, *High-speed noise-free optical quantum memory*, Physical Review A **97**, 042316 (2018).

[29] G. Buser, R. Mottola, B. Cotting, J. Wolters, and P. Treutlein, *Single-photon storage in a ground-state vapor cell quantum memory*, PRX Quantum **3**, 020349 (2022).

[30] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, *Quantum repeaters: The role of imperfect local operations in quantum communication*, Physical Review Letters **81**, 5932 (1998).

[31] S.-K. Liao, W.-Q. Cai, W.-Y. Liu, L. Zhang, Y. Li, J.-G. Ren, J. Yin, Q. Shen, Y. Cao, Z.-P. Li, F.-Z. Li, X.-W. Chen, L.-H. Sun, J.-J. Jia, J.-C. Wu, X.-J. Jiang, J.-F. Wang, Y.-M. Huang, Q. Wang, Y.-L. Zhou, L. Deng, T. Xi, L. Ma, T. Hu, Q. Zhang, Y.-A. Chen, N.-L. Liu, X.-B. Wang, Z.-C. Zhu, C.-Y. Lu, R. Shu, C.-Z. Peng, J.-Y. Wang, and J.-W. Pan, *Satellite-to-ground quantum key distribution*, Nature **549**, 43 (2017).

[32] J.-G. Ren, P. Xu, H.-L. Yong, L. Zhang, S.-K. Liao, J. Yin, W.-Y. Liu, W.-Q. Cai, M. Yang, L. Li, K.-X. Yang, X. Han, Y.-Q. Yao, J. Li, H.-Y. Wu, S. Wan, L. Liu, D.-Q. Liu, Y.-W. Kuang, Z.-P. He, P. Shang, C. Guo, R.-H. Zheng, K. Tian, Z.-C. Zhu, N.-L. Liu, C.-Y. Lu, R. Shu, Y.-A. Chen, C.-Z. Peng, J.-Y. Wang, and J.-W. Pan, *Ground-to-satellite quantum teleportation*, Nature **549**, 70 (2017).

[33] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, *Fundamental limits of repeaterless quantum communications*, Nature Communications **8**, 15043 (2017).

[34] S. Pirandola, *Limits and security of free-space quantum communications*, Physical Review Research **3**, 013279 (2021).

[35] W. K. Wootters and W. H. Zurek, *A single quantum cannot be cloned*, Nature **299**, 802 (1982).

[36] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels*, Physical Review Letters **70**, 1895 (1993).

[37] S. Pirandola, *End-to-end capacities of a quantum communication network*, Communications Physics **2**, 51 (2019).

[38] M. Żukowski, A. Zeilinger, M. A. Horne, and A. K. Ekert, *"Event-ready-detectors" Bell experiment via entanglement swapping*, Physical Review Letters **71**, 4287 (1993).

[39] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, *Long-distance quantum communication with atomic ensembles and linear optics*, Nature **414**, 413 (2001).

[40] C. Simon, H. de Riedmatten, M. Afzelius, N. Sangouard, H. Zbinden, and N. Gisin, *Quantum repeaters with photon pair sources and multimode memories*, Physical Review Letters **98**, 190503 (2007).

[41] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin, *Quantum repeaters based on atomic ensembles and linear optics*, Reviews of Modern Physics **83**, 33 (2011).

[42] A. Kuzmich, W. P. Bowen, A. D. Boozer, A. Boca, C. W. Chou, L.-M. Duan, and H. J. Kimble, *Generation of nonclassical photon pairs for scalable quantum communication with atomic ensembles*, Nature **423**, 731 (2003).

[43] N. Sangouard, C. Simon, J. Minář, H. Zbinden, H. de Riedmatten, and N. Gisin, *Long-distance entanglement distribution with single-photon sources*, Physical Review A **76**, 050301(R) (2007).

[44] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien, *Quantum computers*, Nature **464**, 45 (2010).

[45] E. Knill, R. Laflamme, and G. J. Milburn, *A scheme for efficient quantum computation with linear optics*, Nature **409**, 46 (2001).

[46] D. E. Browne and T. Rudolph, *Resource-efficient linear optical quantum computation*, Physical Review Letters **95**, 010501 (2005).

[47] T. P. Bodiya and L.-M. Duan, *Scalable generation of graph-state entanglement through realistic linear optics*, Physical Review Letters **97**, 143601 (2006).

[48] S. Aaronson and A. Arkhipov, *The computational complexity of linear optics*, in *STOC '11: Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 333–342 (ACM Press) (2011).

[49] J. Nunn, N. K. Langford, W. S. Kolthammer, T. F. M. Champion, M. R. Sprague, P. S. Michelberger, X.-M. Jin, D. G. England, and I. A. Walmsley, *Enhancing multiphoton rates with quantum memories*, Physical Review Letters **110**, 133601 (2013).

[50] S. E. Thomas, *Efficient, low noise, mode-selective quantum memory*, Ph.D. thesis, Imperial College London (2019).

[51] K. F. Reim, J. Nunn, X.-M. Jin, P. S. Michelberger, T. F. M. Champion, D. G. England, K. C. Lee, W. S. Kolthammer, N. K. Langford, and I. A. Walmsley, *Multipulse addressing of a Raman quantum memory: Configurable beam splitting and efficient readout*, Physical Review Letters **108**, 263602 (2012).

[52] J. Kitching, *Chip-scale atomic devices*, Applied Physics Reviews **5**, 031302 (2018).

[53] M. Gündoğan, J. S. Sidhu, V. Henderson, L. Mazzarella, J. Wolters, D. K. L. Oi, and M. Krutzik, *Proposal for space-borne quantum memories for global quantum networking*, npj Quantum Information **7**, 128 (2021).

[54] N. B. Phillips, A. V. Gorshkov, and I. Novikova, *Light storage in an optically thick atomic ensemble under conditions of electromagnetically induced transparency and four-wave mixing*, Physical Review A **83**, 063823 (2011).

[55] P. S. Michelberger, T. F. M. Champion, M. R. Sprague, K. T. Kaczmarek, M. Barbieri, X. M. Jin, D. G. England, W. S. Kolthammer, D. J. Saunders, J. Nunn, and I. A. Walmsley, *Interfacing GHz-bandwidth heralded single photons with a warm vapour Raman memory*, New Journal of Physics **17**, 043006 (2015).

[56] J. Wolters, G. Buser, A. Horsley, L. Béguin, A. Jöckel, J.-P. Jahn, R. J. Warburton, and P. Treutlein, *Simple atomic quantum memory suitable for semiconductor quantum dot single photons*, Physical Review Letters **119**, 060502 (2017).

[57] K. F. Reim, P. Michelberger, K. C. Lee, J. Nunn, N. K. Langford, and I. A. Walmsley, *Single-photon-level quantum memory at room temperature*, Physical Review Letters **107**, 053603 (2011).

[58] D. J. Saunders, J. H. D. Munns, T. F. M. Champion, C. Qiu, K. T. Kaczmarek, E. Poem, P. M. Ledingham, I. A. Walmsley, and J. Nunn, *Cavity-enhanced room-temperature broadband Raman memory*, Physical Review Letters **116**, 090501 (2016).

[59] O. Katz and O. Firstenberg, *Light storage for one second in room-temperature alkali vapor*, Nature Communications **9**, 2074 (2018).

[60] P. Senellart, G. Solomon, and A. White, *High-performance semiconductor quantum-dot single-photon sources*, Nature Nanotechnology **12**, 1026 (2017).

[61] L. Zhai, M. C. Löbl, G. N. Nguyen, J. Ritzmann, A. Javadi, C. Spinnler, A. D. Wieck, A. Ludwig, and R. J. Warburton, *Low-noise GaAs quantum dots for quantum photonics*, Nature Communications **11**, 4745 (2020).

[62] L. Zhai, G. N. Nguyen, C. Spinnler, J. Ritzmann, M. C. Löbl, A. D. Wieck, A. Ludwig, A. Javadi, and R. J. Warburton, *Quantum interference of identical photons from remote GaAs quantum dots*, Nature Nanotechnology **17**, 829 (2022).

[63] N. Tomm, A. Javadi, N. O. Antoniadis, D. Najer, M. C. Löbl, A. R. Korsch, R. Schott, S. R. Valentin, A. D. Wieck, A. Ludwig, and R. J. Warburton, *A bright and fast source of coherent single photons*, Nature Nanotechnology **16**, 399 (2021).

[64] R. Mottola, G. Buser, C. Müller, T. Kroh, A. Ahlrichs, S. Ramelow, O. Benson, P. Treutlein, and J. Wolters, *An efficient, tunable, and robust source of narrowband photon pairs at the ^{87}Rb D1 line*, Optics Express **28**, 3159 (2020).

[65] R. Mottola, G. Buser, and P. Treutlein, *Optical memory in a microfabricated rubidium vapor cell*, Physical Review Letters **131**, 260801 (2023).

[66] R. Mottola, G. Buser, and P. Treutlein, *Electromagnetically induced transparency and optical pumping in the hyperfine Paschen-Back regime*, Physical Review A **108**, 062820 (2023).

[67] I. Aharonovich, D. Englund, and M. Toth, *Solid-state single-photon emitters*, Nature Photonics **10**, 631 (2016).

[68] T. Gaebel, I. Popa, A. Gruber, M. Domhan, F. Jelezko, and J. Wrachtrup, *Stable single-photon source in the near infrared*, New Journal of Physics **6**, 98 (2004).

[69] E. Wu, J. R. Rabeau, G. Roger, F. Treussart, H. Zeng, P. Grangier, S. Prawer, and J.-F. Roch, *Room temperature triggered single-photon source in the near infrared*, New Journal of Physics **9**, 434 (2007).

[70] M. Hennrich, T. Legero, A. Kuhn, and G. Rempe, *Photon statistics of a non-stationary periodically driven single-photon source*, New Journal of Physics **6**, 86 (2004).

[71] C. Maurer, C. Becher, C. Russo, J. Eschner, and R. Blatt, *A single-photon source based on a single Ca^+ ion*, New Journal of Physics **6**, 94 (2004).

[72] M. Steiner, A. Hartschuh, R. Korlacki, and A. J. Meixner, *Highly efficient, tunable single photon source based on single molecules*, Applied Physics Letters **90**, 183122 (2007).

[73] C. Zhang, Y. Huang, B. Liu, C. Li, and G. Guo, *Spontaneous parametric down-conversion sources for multiphoton experiments*, Advanced Quantum Technologies **4**, 2000132 (2021).

[74] M. Rambach, A. Nikolova, T. J. Weinhold, and A. G. White, *Sub-megahertz linewidth single photon source*, APL Photonics **1**, 096101 (2016).

[75] J. Fekete, D. Rieländer, M. Cristiani, and H. de Riedmatten, *Ultranarrow-band photon-pair source compatible with solid state quantum memories and telecommunication networks*, Physical Review Letters **110**, 220502 (2013).

[76] A. Ahlrichs and O. Benson, *Bright source of indistinguishable photons based on cavity-enhanced parametric down-conversion utilizing the cluster effect*, Applied Physics Letters **108**, 021111 (2016).

[77] A. B. U'Ren, C. Silberhorn, K. Banaszek, and I. A. Walmsley, *Efficient conditional preparation of high-fidelity single photon states for fiber-optic quantum networks*, Physical Review Letters **93**, 093601 (2004).

[78] J. Zhao, C. Ma, M. Rüsing, and S. Mookherjea, *High quality entangled photon pair generation in periodically poled thin-film lithium niobate waveguides*, Physical Review Letters **124**, 163603 (2020).

[79] M. Fiorentino, P. Voss, J. Sharping, and P. Kumar, *All-fiber photon-pair source for quantum communications*, IEEE Photonics Technology Letters **14**, 983 (2002).

[80] J. G. Rarity, J. Fulconis, J. Duligall, W. J. Wadsworth, and P. S. J. Russell, *Photonic crystal fiber source of correlated photon pairs*, Optics Express **13**, 534 (2005).

[81] J. Fan, A. Migdall, and L. J. Wang, *Efficient generation of correlated photon pairs in a microstructure fiber*, Optics Letters **30**, 3368 (2005).

[82] J. E. Sharping, K. F. Lee, M. A. Foster, A. C. Turner, B. S. Schmidt, M. Lipson, A. L. Gaeta, and P. Kumar, *Generation of correlated photons in nanoscale silicon waveguides*, Optics Express **14**, 12388 (2006).

[83] K.-i. Harada, H. Takesue, H. Fukuda, T. Tsuchizawa, T. Watanabe, K. Yamada, Y. Tokura, and S.-i. Itabashi, *Generation of high-purity entangled photon pairs using silicon wire waveguide*, Optics Express **16**, 20368 (2008).

[84] H. Takesue, Y. Tokura, H. Fukuda, T. Tsuchizawa, T. Watanabe, K. Yamada, and S.-i. Itabashi, *Entanglement generation using silicon wire waveguide*, Applied Physics Letters **91**, 201108 (2007).

[85] B. Srivathsan, G. K. Gulati, B. Chng, G. Maslennikov, D. Matsukevich, and C. Kurtsiefer, *Narrow band source of transform-limited photon pairs via four-wave mixing in a cold atomic ensemble*, Physical Review Letters **111**, 123602 (2013).

[86] O. Davidson, R. Finkelstein, E. Poem, and O. Firstenberg, *Bright multiplexed source of indistinguishable single photons with tunable GHz-bandwidth at room temperature*, New Journal of Physics **23**, 073050 (2021).

[87] A. I. Lvovsky, H. Hansen, T. Aichele, O. Benson, J. Mlynek, and S. Schiller, *Quantum state reconstruction of the single-photon Fock state*, Physical Review Letters **87**, 050402 (2001).

[88] R. Loudon, *The quantum theory of light* (OUP Oxford) (2000).

[89] C. C. Gerry and P. Knight, *Introductory quantum optics*, (Cambridge University Press, Cambridge) (2004).

[90] C. Müller, A. Ahlrichs, and O. Benson, *General and complete description of temporal photon correlations in cavity-enhanced spontaneous parametric down-conversion*, Physical Review A **102**, 053504 (2020).

[91] C. K. Hong, Z. Y. Ou, and L. Mandel, *Measurement of subpicosecond time intervals between two photons by interference*, Physical Review Letters **59**, 2044 (1987).

[92] F. Bouchard, A. Sit, Y. Zhang, R. Fickler, F. M. Miatto, Y. Yao, F. Sciarrino, and E. Karimi, *Two-photon interference: The Hong-Ou-Mandel effect*, Reports on Progress in Physics **84**, 012402 (2021).

[93] Z.-Y. J. Ou, *Multi-photon quantum interference*, (Springer New York, NY) (2007).

[94] P. Jobez, C. Laplane, N. Timoney, N. Gisin, A. Ferrier, P. Goldner, and M. Afzelius, *Coherent spin control at the quantum level in an ensemble-based optical memory*, Physical Review Letters **114**, 230502 (2015).

[95] A. B. U'Ren, Y. Jeronimo-Moreno, and H. Garcia-Gracia, *Generation of Fourier-transform-limited heralded single photons*, Physical Review A **75**, 023810 (2007).

[96] Y.-J. Wei, Y.-M. He, M.-C. Chen, Y.-N. Hu, Y. He, D. Wu, C. Schneider, M. Kamp, S. Höfling, C.-Y. Lu, and J.-W. Pan, *Deterministic and robust generation of single photons from a single quantum dot with* 99.5 % *indistinguishability using adiabatic rapid passage*, Nano Letters **14**, 6515 (2014).

[97] P. Lodahl, A. Ludwig, and R. J. Warburton, *A deterministic source of single photons*, Physics Today **75**, 44 (2022).

[98] J.-P. Jahn, M. Munsch, L. Béguin, A. V. Kuhlmann, M. Renggli, Y. Huo, F. Ding, R. Trotta, M. Reindl, O. G. Schmidt, A. Rastelli, P. Treutlein, and R. J. Warburton, *An artificial Rb atom in a semiconductor with lifetime-limited linewidth*, Physical Review B **92**, 245439 (2015).

[99] L. Zhai, M. C. Löbl, J.-P. Jahn, Y. Huo, P. Treutlein, O. G. Schmidt, A. Rastelli, and R. J. Warburton, *Large-range frequency tuning of a narrow-linewidth quantum emitter*, Applied Physics Letters **117**, 083106 (2020).

[100] S. Ates, I. Agha, A. Gulinatti, I. Rech, A. Badolato, and K. Srinivasan, *Improving the performance of bright quantum dot single photon sources using temporal filtering via amplitude modulation*, Scientific Reports **3**, 1397 (2013).

[101] L. Béguin, J.-P. Jahn, J. Wolters, M. Reindl, Y. Huo, R. Trotta, A. Rastelli, F. Ding, O. G. Schmidt, P. Treutlein, and R. J. Warburton, *On-demand semiconductor source of 780-nm single photons with controlled temporal wave packets*, Physical Review B **97**, 205304 (2018).

[102] L. Shen, J. Lee, A. W. Hartanto, P. Tan, and C. Kurtsiefer, *Wide-range wavelength-tunable photon-pair source for characterizing single-photon detectors*, Optics Express **29**, 3415 (2021).

[103] R. Boyd, *Nonlinear optics (third edition)* (Academic Press) (2008).

[104] M. Yamada, N. Nada, M. Saitoh, and K. Watanabe, *First-order quasi-phase matched LiNbO3 waveguide periodically poled by applying an external field for efficient blue second-harmonic generation*, Applied Physics Letters **62**, 435 (1993).

[105] J. Chen, A. J. Pearlman, A. Ling, J. Fan, and A. L. Migdall, *A versatile waveguide source of photon pairs for chip-scale quantum information processing*, Optics Express **17**, 6727 (2009).

[106] J.-y. Chen, Y. Meng Sua, Z.-h. Ma, C. Tang, Z. Li, and Y.-p. Huang, *Efficient parametric frequency conversion in lithium niobate nanophotonic chips*, OSA Continuum **2**, 2914 (2019).

[107] H.-S. Zhong, Y. Li, W. Li, L.-C. Peng, Z.-E. Su, Y. Hu, Y.-M. He, X. Ding, W. Zhang, H. Li, L. Zhang, Z. Wang, L. You, X.-L. Wang, X. Jiang, L. Li, Y.-A. Chen, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, *12-photon entanglement and scalable scattershot boson sampling with optimal entangled-photon pairs from parametric down-conversion*, Physical Review Letters **121**, 250505 (2018).

[108] A. Ahlrichs, *Triply-resonant cavity-enhanced spontaneous parametric down-conversion*, Ph.D. thesis, Humboldt-Universität zu Berlin (2019).

[109] C. S. Chuu, G. Y. Yin, and S. E. Harris, *A miniature ultrabright source of temporally long, narrowband biphotons*, Applied Physics Letters **101**, 051108 (2012).

[110] K.-H. Luo, H. Herrmann, S. Krapick, B. Brecht, R. Ricken, V. Quiring, H. Suche, W. Sohler, and C. Silberhorn, *Direct generation of genuine single-longitudinal-mode narrowband photon pairs*, New Journal of Physics **17**, 073039 (2015).

[111] P. Kok and B. W. Lovett, *Introduction to optical quantum information processing*, (Springer New York, NY) (2010).

[112] M. Guo, S. Liu, W. Sun, M. Ren, F. Wang, and M. Zhong, *Rare-earth quantum memories: The experimental status quo*, Frontiers of Physics **18**, 21303 (2023).

[113] A. L. Alexander, J. J. Longdell, M. J. Sellars, and N. B. Manson, *Photon echoes produced by switching electric fields*, Physical Review Letters **96**, 043602 (2006).

[114] H. de Riedmatten, M. Afzelius, M. U. Staudt, C. Simon, and N. Gisin, *A solid-state light-matter interface at the single-photon level*, Nature **456**, 773 (2008).

[115] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin, *Multimode quantum memory based on atomic frequency combs*, Physical Review A **79**, 052329 (2009).

[116] G. D. Fuchs, G. Burkard, P. V. Klimov, and D. D. Awschalom, *A quantum memory intrinsic to single nitrogen-vacancy centres in diamond*, Nature Physics **7**, 789 (2011).

[117] T. Chanelière, D. N. Matsukevich, S. D. Jenkins, S. Y. Lan, T. A. Kennedy, and A. Kuzmich, *Storage and retrieval of single photons transmitted between remote quantum memories*, Nature **438**, 833 (2005).

[118] K. S. Choi, H. Deng, J. Laurat, and H. J. Kimble, *Mapping photonic entanglement into and out of a quantum memory*, Nature **452**, 67 (2008).

[119] X.-H. Bao, A. Reingruber, P. Dietrich, J. Rui, A. Dück, T. Strassel, L. Li, N.-L. Liu, B. Zhao, and J.-W. Pan, *Efficient and long-lived quantum memory with cold atoms inside a ring cavity*, Nature Physics **8**, 517 (2012).

[120] Y. W. Cho, G. T. Campbell, J. L. Everett, J. Bernu, D. B. Higginbottom, M. T. Cao, J. Geng, N. P. Robins, P. K. Lam, and B. C. Buchler, *Highly efficient optical quantum memory with long coherence time in cold atoms*, Optica **3**, 100 (2016).

[121] P. Vernaz-Gris, K. Huang, M. Cao, A. S. Sheremet, and J. Laurat, *Highly-efficient quantum memory for polarization qubits in a spatially-multiplexed cold atomic ensemble*, Nature Communications **9**, 363 (2018).

[122] H. P. Specht, C. Nölleke, A. Reiserer, M. Uphoff, E. Figueroa, S. Ritter, and G. Rempe, *A single-atom quantum memory*, Nature **473**, 190 (2011).

[123] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, *An elementary quantum network of single atoms in optical cavities*, Nature **484**, 195 (2012).

[124] M. Namazi, C. Kupchak, B. Jordaan, R. Shahrokhshahi, and E. Figueroa, *Ultralow-noise room-temperature quantum memory for polarization qubits*, Physical Review Applied **8**, 034023 (2017).

[125] M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler, *Unconditional room-temperature quantum memory*, Nature Physics **7**, 794 (2011).

[126] R. Finkelstein, E. Poem, O. Michel, O. Lahad, and O. Firstenberg, *Fast, noise-free memory for photon synchronization at room temperature*, Science Advances **4**, eaap8598 (2018).

[127] M. Nilsson and S. Kröll, *Solid state quantum memory using complete absorption and re-emission of photons by tailored and externally controlled inhomogeneous absorption profiles*, Optics Communications **247**, 393 (2005).

[128] B. Kraus, W. Tittel, N. Gisin, M. Nilsson, S. Kröll, and J. I. Cirac, *Quantum memory for nonstationary light fields based on controlled reversible inhomogeneous broadening*, Physical Review A **73**, 020302(R) (2006).

[129] M. D. Lukin, *Colloquium: Trapping and manipulating photon states in atomic ensembles*, Reviews of Modern Physics **75**, 457 (2003).

[130] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in Λ-type optically dense atomic media. II. free-space model*, Physical Review A **76**, 033805 (2007).

[131] A. V. Gorshkov, A. André, M. Fleischhauer, A. S. Sørensen, and M. D. Lukin, *Universal approach to optimal photon storage in atomic media*, Physical Review Letters **98**, 123601 (2007).

[132] J. Khurgin and R. Tucker (editors), *Slow light: Science and applications*, (CRC Press, Boca Raton) (2009).

[133] F. Kaneda, F. Xu, J. Chapman, and P. G. Kwiat, *Quantum-memory-assisted multi-photon generation for efficient quantum information processing*, Optica **4**, 1034 (2017).

[134] K. Shinbrough, D. R. Pearson, B. Fang, E. A. Goldschmidt, and V. O. Lorenz, *Chapter Five - Broadband quantum memory in atomic ensembles*, in S. F. Y. Louis F. DiMauro, Hélène Perrin (editor), *Advances in Atomic, Molecular, and Optical Physics*, volume 72, pp. 297–360 (Elsevier) (2023).

[135] J. L. Dodd and M. A. Nielsen, *Simple operational interpretation of the fidelity of mixed states*, Physical Review A **66** (2002).

[136] Y.-H. Chen, M.-J. Lee, I.-C. Wang, S. Du, Y.-F. Chen, Y.-C. Chen, and I. A. Yu, *Coherent optical memory with high storage efficiency and large fractional delay*, Physical Review Letters **110**, 083601 (2013).

[137] M. Gündoğan, P. M. Ledingham, K. Kutluer, M. Mazzera, and H. de Riedmatten, *Solid state spin-wave quantum memory for time-bin qubits*, Physical Review Letters **114**, 230501 (2015).

[138] E. A. Goldschmidt, F. Piacentini, I. R. Berchera, S. V. Polyakov, S. Peters, S. Kück, G. Brida, I. P. Degiovanni, A. Migdall, and M. Genovese, *Mode reconstruction of a light field by multiphoton statistics*, Physical Review A **88**, 013822 (2013).

[139] Y.-F. Pu, N. Jiang, W. Chang, H.-X. Yang, C. Li, and L.-M. Duan, *Experimental realization of a multiplexed quantum memory with 225 individually accessible memory cells*, Nature Communications **8**, 15359 (2017).

[140] L. Meßner, E. Robertson, L. Esguerra, K. Lüdge, and J. Wolters, *Multiplexed random-access optical memory in warm cesium vapor*, Optics Express **31**, 10150 (2023).

[141] D. V. Vasilyev, I. V. Sokolov, and E. S. Polzik, *Quantum memory for images: A quantum hologram*, Physical Review A **77**, 020302(R) (2008).

[142] M. Businger, L. Nicolas, T. S. Mejia, A. Ferrier, P. Goldner, and M. Afzelius, *Non-classical correlations over 1250 modes between telecom photons and 979-nm photons stored in* $^{171}Yb^{3+}{:}Y_2SiO_5$, Nature Communications **13** (2022).

[143] P. Kumar, *Quantum frequency conversion*, Optics Letters **15**, 1476 (1990).

[144] J. H. Weber, B. Kambs, J. Kettler, S. Kern, J. Maisch, H. Vural, M. Jetter, S. L. Portalupi, C. Becher, and P. Michler, *Two-photon interference in the telecom C-band after frequency conversion of photons from remote quantum emitters*, Nature Nanotechnology **14**, 23 (2018).

[145] T. van Leent, M. Bock, R. Garthoff, K. Redeker, W. Zhang, T. Bauer, W. Rosenfeld, C. Becher, and H. Weinfurter, *Long-distance distribution of atom-photon entanglement at telecom wavelength*, Physical Review Letters **124**, 010510 (2020).

[146] M. Hosseini, B. Sparkes, G. Campbell, P. Lam, and B. Buchler, *High efficiency coherent optical memory with warm rubidium vapour*, Nature Communications **2**, 174 (2011).

[147] D. F. Phillips, A. Fleischhauer, A. Mair, R. L. Walsworth, and M. D. Lukin, *Storage of light in atomic vapor*, Physical Review Letters **86**, 783 (2001).

[148] N. B. Phillips, A. V. Gorshkov, and I. Novikova, *Optimal light storage in atomic vapor*, Physical Review A **78**, 023801 (2008).

[149] K. F. Reim, J. Nunn, V. O. Lorenz, B. J. Sussman, K. C. Lee, N. K. Langford, D. Jaksch, and I. A. Walmsley, *Towards high-speed optical quantum memories*, Nature Photonics **4**, 218 (2010).

[150] S. E. Harris, *Electromagnetically induced transparency*, Physics Today **50**, 36 (1997).

[151] M. Fleischhauer, A. Imamoglu, and J. P. Marangos, *Electromagnetically induced transparency: Optics in coherent media*, Reviews of Modern Physics **77**, 633 (2005).

[152] R. Finkelstein, S. Bali, O. Firstenberg, and I. Novikova, *A practical guide to electromagnetically induced transparency in atomic vapor*, New Journal of Physics **25**, 035001 (2023).

[153] S. E. Harris, J. E. Field, and A. Kasapi, *Dispersive properties of electromagnetically induced transparency*, Physical Review A **46**, R29 (1992).

[154] L. V. Hau, S. E. Harris, Z. Dutton, and C. H. Behroozi, *Light speed reduction to 17 metres per second in an ultracold atomic gas*, Nature **397**, 594 (1999).

[155] M. Fleischhauer, S. F. Yelin, and M. D. Lukin, *How to trap photons? Storing single-photon quantum states in collective atomic excitations*, Optics Communications **179**, 395 (2000).

[156] M. Fleischhauer and M. D. Lukin, *Quantum memory for photons: Dark-state polaritons*, Physical Review A **65**, 022314 (2002).

[157] C. Mewes and M. Fleischhauer, *Decoherence in collective quantum memories for photons*, Physical Review A **72**, 022327 (2005).

[158] I. Novikova, R. L. Walsworth, and Y. Xiao, *Electromagnetically induced transparency-based slow and stored light in warm atoms*, Laser & Photonics Reviews **6**, 333 (2011).

[159] J. Nunn, *Quantum memory in atomic ensembles*, Ph.D. thesis, St. John's College, University of Oxford (2008).

[160] S. Manz, T. Fernholz, J. Schmiedmayer, and J.-W. Pan, *Collisional decoherence during writing and reading quantum states*, Physical Review A **75**, 040101(R) (2007).

[161] D. L. Rousseau, G. D. Patterson, and P. F. Williams, *Resonance Raman scattering and collision-induced redistribution scattering in I_2*, Physical Review Letters **34**, 1306 (1975).

[162] K.-I. Harada, T. Kanbashi, M. Mitsunaga, and K. Motomura, *Competition between electromagnetically induced transparency and stimulated Raman scattering*, Physical Review A **73**, 013807 (2006).

[163] G. S. Agarwal, T. N. Dey, and D. J. Gauthier, *Competition between electromagnetically induced transparency and Raman processes*, Physical Review A **74**, 043805 (2006).

[164] N. Lauk, C. O'Brien, and M. Fleischhauer, *Fidelity of photon propagation in electromagnetically induced transparency in the presence of four-wave mixing*, Physical Review A **88**, 013823 (2013).

[165] P. R. S. Carvalho, L. E. E. de Araujo, and J. W. R. Tabosa, *Angular dependence of an electromagnetically induced transparency resonance in a Doppler-broadened atomic vapor*, Physical Review A **70**, 063818 (2004).

[166] B. Zhao, Y.-A. Chen, X.-H. Bao, T. Strassel, C.-S. Chuu, X.-M. Jin, J. Schmiedmayer, Z.-S. Yuan, S. Chen, and J.-W. Pan, *A millisecond quantum memory for scalable quantum networks*, Nature Physics **5**, 95 (2008).

[167] B. Gouraud, D. Maxein, A. Nicolas, O. Morin, and J. Laurat, *Demonstration of a memory for tightly guided light in an optical nanofiber*, Physical Review Letters **114**, 180503 (2015).

[168] A. Horsley, *High resolution field imaging with atomic vapor cells*, Ph.D. thesis, Universität Basel (2015).

[169] J. C. Allred, R. N. Lyman, T. W. Kornack, and M. V. Romalis, *High-sensitivity atomic magnetometer unaffected by spin-exchange relaxation*, Physical Review Letters **89**, 130801 (2002).

[170] S. Kadlecek, L. W. Anderson, and T. G. Walker, *Field dependence of spin relaxation in a dense Rb vapor*, Physical Review Letters **80**, 5512 (1998).

[171] O. Davidson, O. Yogev, E. Poem, and O. Firstenberg, *Fast, noise-free atomic optical memory with 35-percent end-to-end efficiency*, Communications Physics **6**, 631 (2023).

[172] R. Finkelstein, O. Lahad, I. Cohen, O. Davidson, S. Kiriati, E. Poem, and O. Firstenberg, *Continuous protection of a collective state from inhomogeneous dephasing*, Physical Review X **11**, 011008 (2021).

[173] S. Gao, O. Lazo-Arjona, B. Brecht, K. T. Kaczmarek, S. E. Thomas, J. Nunn, P. M. Ledingham, D. J. Saunders, and I. A. Walmsley, *Optimal coherent filtering for single noisy photons*, Physical Review Letters **123**, 213604 (2019).

[174] S. Thomas, S. Sagona-Stophel, Z. Schofield, I. Walmsley, and P. Ledingham, *Single-photon-compatible telecommunications-band quantum memory in a hot atomic gas*, Physical Review Applied **19**, L031005 (2023).

[175] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in Λ-type optically dense atomic media. I. cavity model*, Physical Review A **76**, 033804 (2007).

[176] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in Λ-type optically dense atomic media. III. effects of inhomogeneous broadening*, Physical Review A **76**, 033806 (2007).

[177] A. V. Gorshkov, T. Calarco, M. D. Lukin, and A. S. Sørensen, *Photon storage in Λ-type optically dense atomic media. IV. optimal control using gradient ascent*, Physical Review A **77**, 043806 (2008).

[178] M. T. Rakher, R. J. Warburton, and P. Treutlein, *Prospects for storage and retrieval of a quantum-dot single photon in an ultracold ^{87}Rb ensemble*, Physical Review A **88**, 053834 (2013).

[179] K. Surmacz, J. Nunn, K. Reim, K. C. Lee, V. O. Lorenz, B. Sussman, I. A. Walmsley, and D. Jaksch, *Efficient spatially resolved multimode quantum memory*, Physical Review A **78**, 033806 (2008).

[180] L. N. Trefethen, *Spectral methods in MATLAB* (Society for Industrial and Applied Mathematics, Philadelphia) (2000).

[181] J. P. Boyd, *Chebyshev and Fourier spectral methods* (DOVER Publications, Inc.) (2001).

[182] B. Cotting, *Spontaneous parametric down-conversion heralded single-photon source for quantum memory applications*, Master's thesis, University of Basel, École polytechnique fédérale de Lausanne (2021).

[183] A. Vanselow, P. Kaufmann, H. M. Chrzanowski, and S. Ramelow, *Ultra-broadband SPDC for spectrally far separated photon pairs*, Optics Letters **44**, 4638 (2019).

[184] M. Mehmet, H. Vahlbruch, N. Lastzka, K. Danzmann, and R. Schnabel, *Observation of squeezed states with strong photon-number oscillations*, Physical Review A **81**, 013814 (2010).

[185] G. Breitenbach, T. Müller, S. F. Pereira, J.-P. Poizat, S. Schiller, and J. Mlynek, *Squeezed vacuum from a monolithic optical parametric oscillator*, Journal of the Optical Society of America B **12**, 2304 (1995).

[186] A. Brieussel, Y. Shen, G. Campbell, G. Guccione, J. Janousek, B. Hage, B. C. Buchler, N. Treps, C. Fabre, F. Z. Fang, X. Y. Li, T. Symul, and P. K. Lam, *Squeezed light from a diamond-turned monolithic cavity*, Optics Express **24**, 4042 (2016).

[187] S. Schiller, R. Bruckmeier, M. Schalke, K. Schneider, and J. Mlynek, *Quantum nondemolition measurements and generation of individually squeezed twin beams by a degenerate optical parametric amplifier*, Europhysics Letters **36**, 361 (1996).

[188] M. Mehmet, S. Ast, T. Eberle, S. Steinlechner, H. Vahlbruch, and R. Schnabel, *Squeezed light at 1550 nm with a quantum noise reduction of 12.3 dB*, Optics Express **19**, 25763 (2011).

[189] R. Guo, X. Jia, C. Xie, and K. Peng, *A portable multi-purpose non-classical light source*, Optics Communications **211**, 243 (2002).

[190] Y. Zhou, X. Jia, F. Li, C. Xie, and K. Peng, *Experimental generation of 8.4 dB entangled state with an optical cavity involving a wedged type-II nonlinear crystal*, Optics Express **23**, 4952 (2015).

[191] E. Pomarico, B. Sanguinetti, N. Gisin, R. Thew, H. Zbinden, G. Schreiber, A. Thomas, and W. Sohler, *Waveguide-based OPO source of entangled photon pairs*, New Journal of Physics **11**, 113042 (2009).

[192] E. Pomarico, B. Sanguinetti, C. I. Osorio, H. Herrmann, and R. T. Thew, *Engineering integrated pure narrow-band photon sources*, New Journal of Physics **14** (2012).

[193] A. Siegman, *Lasers* (University Science Books) (1986).

[194] S. Emanueli and A. Arie, *Temperature-dependent dispersion equations for $KTiOPO_4$ and $KTiOAsO_4$*, Applied Optics **42**, 6661 (2003).

[195] J. Steinlechner, S. Ast, C. Krüger, A. Singh, T. Eberle, V. Händchen, and R. Schnabel, *Absorption measurements of periodically poled potassium titanyl phosphate (PPKTP) at* 775 nm *and* 1550 nm, Sensors **13**, 565 (2013).

[196] M. Scholz, L. Koch, and O. Benson, *Analytical treatment of spectral properties and signal–idler intensity correlations for a double-resonant optical parametric oscillator far below threshold*, Optics Communications **282**, 3518 (2009).

[197] A. Araya, S. Telada, K. Tochikubo, S. Taniguchi, R. Takahashi, K. Kawabe, D. Tatsumi, T. Yamazaki, S. Kawamura, S. Miyoki, S. Moriwaki, M. Musha, S. Nagano, M.-K. Fujimoto, K. Horikoshi, N. Mio, Y. Naito, A. Takamori, and K. Yamamoto, *Absolute-length determination of a long-baseline Fabry-Perot cavity by means of resonating modulation sidebands*, Applied Optics **38**, 2848 (1999).

[198] B. Boulanger, M. M. Fejer, R. Blachman, and P. F. Bordui, *Study of $KTiOPO_4$ gray-tracking at 1064, 532, and 355 nm*, Applied Physics Letters **65**, 2401 (1994).

[199] X. Mu and Y. J. Ding, *Investigation of damage mechanisms of $KTiOPO_4$ crystals by use of a continuous-wave argon laser*, Applied Optics **39**, 3099 (2000).

[200] J. K. Tyminski, *Photorefractive damage in KTP used as second-harmonic generator*, Journal of Applied Physics **70**, 5570 (1991).

[201] G. M. Loiacono, D. N. Loiacono, T. McGee, and M. Babb, *Laser damage formation in $KTiOPO_4$ and $KTiOAsO_4$ crystals: Grey tracks*, Journal of Applied Physics **72**, 2705 (1992).

[202] R. Blachman, P. F. Bordui, and M. M. Fejer, *Laser-induced photochromic damage in potassium titanyl phosphate*, Applied Physics Letters **64**, 1318 (1994).

[203] P. Sekatski, N. Sangouard, F. Bussières, C. Clausen, N. Gisin, and H. Zbinden, *Detector imperfections in photon-pair source characterization*, Journal of Physics B: Atomic, Molecular and Optical Physics **45**, 124016 (2012).

[204] U. Herzog, M. Scholz, and O. Benson, *Theory of biphoton generation in a single-resonant optical parametric oscillator far below threshold*, Physical Review A **77**, 023826 (2008).

[205] A. E. B. Nielsen and K. Mølmer, *Multimode analysis of the light emitted from a pulsed optical parametric oscillator*, Physical Review A **76**, 033832 (2007).

[206] D. Rieländer, A. Lenhard, M. Mazzera, and H. de Riedmatten, *Cavity enhanced telecom heralded single photons for spin-wave solid state quantum memories*, New Journal of Physics **18**, 123013 (2016).

[207] A. Christ, K. Laiho, A. Eckstein, K. N. Cassemiro, and C. Silberhorn, *Probing multimode squeezing with correlation functions*, New Journal of Physics **13**, 033027 (2011).

[208] B. E. A. Saleh and M. C. Teich, *Fundamentals of photonics*, (John Wiley & Sons, Inc.) (1991).

[209] C. W. Chou, S. V. Polyakov, A. Kuzmich, and H. J. Kimble, *Single-photon generation from stored excitation in an atomic ensemble*, Physical Review Letters **92**, 213601 (2004).

[210] S. H. Wemple and M. DiDomenico, *Theory of the elasto-optic effect in nonmetallic crystals*, Physical Review B **1**, 193 (1970).

[211] J. A. Zielińska, A. Zukauskas, C. Canalias, M. A. Noyan, and M. W. Mitchell, *Fully-resonant, tunable, monolithic frequency conversion as a coherent UVA source*, Optics Express **25**, 1142 (2017).

[212] C. Groves, C. H. Tan, J. P. R. David, G. J. Rees, and M. M. Hayat, *Exponential time response in analogue and Geiger mode avalanche photodiodes*, IEEE Transactions on Electron Devices **52**, 1527 (2005).

[213] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, *Laser phase and frequency stabilization using an optical resonator*, Applied Physics B **31**, 97 (1983).

[214] E. D. Black, *An introduction to Pound-Drever-Hall laser frequency stabilization*, American Journal of Physics **69**, 79 (2001).

[215] J. I. Thorpe, K. Numata, and J. Livas, *Laser frequency stabilization and control through offset sideband locking to optical cavities*, Optics Express **16**, 15980 (2008).

[216] M. Scholz, L. Koch, R. Ullmann, and O. Benson, *Single-mode operation of a high-brightness narrow-band single-photon source*, Applied Physics Letters **94**, 20 (2009).

[217] D. A. Steck, *Rubidium 87 D line data*, available at https://steck.us/alkalidata (revision 2.2.2, 2021).

[218] A. Corney, *Atomic and laser spectroscopy*, (Oxford University Press) (2006).

[219] T. F. Gallagher, *Rydberg atoms*, Cambridge Monographs on Atomic, Molecular and Chemical Physics (Cambridge University Press) (1994).

[220] L. Weller, K. S. Kleinbach, M. A. Zentile, S. Knappe, C. S. Adams, and I. G. Hughes, *Absolute absorption and dispersion of a rubidium vapour in the hyperfine Paschen-Back regime*, Journal of Physics B: Atomic, Molecular and Optical Physics **45**, 215005 (2012).

[221] B. A. Olsen, B. Patton, Y.-Y. Jau, and W. Happer, *Optical pumping and spectroscopy of Cs vapor at high magnetic field*, Physical Review A **84**, 063410 (2011).

[222] A. Sargsyan, G. Hakhumyan, R. Mirzoyan, and D. Sarkisyan, *Investigation of atomic transitions of cesium in strong magnetic fields by an optical half-wavelength cell*, JETP Letters **98**, 441 (2013).

[223] A. Sargsyan, G. Hakhumyan, C. Leroy, Y. Pashayan-Leroy, A. Papoyan, and D. Sarkisyan, *Hyperfine Paschen-Back regime realized in Rb nanocell*, Optics Letters **37**, 1379 (2012).

[224] A. Sargsyan, A. Tonoyan, G. Hakhumyan, C. Leroy, Y. Pashayan-Leroy, and D. Sarkisyan, *Complete hyperfine Paschen-Back regime at relatively small magnetic fields realized in potassium nano-cell*, EPL **110** (2015).

[225] C. Umfer, L. Windholz, and M. Musso, *Investigations of the sodium and lithium D-lines in strong magnetic fields*, Zeitschrift für Physik D - Atoms, Molecules and Clusters **25**, 23 (1992).

[226] D. R. Lide (editor), *CRC handbook of chemistry and physics*, 82nd ed. (CRC press, Boca Raton) (2001).

[227] G. Breit and I. I. Rabi, *Measurement of nuclear spin*, Physical Review **38**, 2082 (1931).

[228] P. Siddons, C. S. Adams, C. Ge, and I. G. Hughes, *Absolute absorption on rubidium D lines: Comparison between theory and experiment*, Journal of Physics B: Atomic, Molecular and Optical Physics **41**, 155004 (2008).

[229] M. A. Zentile, J. Keaveney, L. Weller, D. J. Whiting, C. S. Adams, and I. G. Hughes, *ElecSus: A program to calculate the electric susceptibility of an atomic ensemble*, Computer Physics Communications **189**, 162 (2015).

[230] D. Pizzey, J. D. Briscoe, F. D. Logue, F. S. Ponciano-Ojeda, S. A. Wrathmall, and I. G. Hughes, *Laser spectroscopy of hot atomic vapours: From 'scope to theoretical fit*, New Journal of Physics **24**, 125001 (2022).

[231] J. Keaveney, C. S. Adams, and I. G. Hughes, *ElecSus: Extension to arbitrary geometry magneto-optics*, Computer Physics Communications **224**, 311 (2018).

[232] A. Sargsyan, G. Hakhumyan, C. Leroy, Y. Pashayan-Leroy, A. Papoyan, D. Sarkisyan, and M. Auzinsh, *Hyperfine Paschen-Back regime in alkali metal atoms: Consistency of two theoretical considerations and experiment*, Journal of the Optical Society of America B **31**, 1046 (2014).

[233] W. Happer, *Optical pumping*, Reviews of Modern Physics **44**, 169 (1972).

[234] M. Pellaton, C. Affolderbach, Y. Pétremand, N. de Rooij, and G. Mileti, *Study of laser-pumped double-resonance clock signals using a microfabricated cell*, Physica Scripta **T149**, 014013 (2012).

[235] A. Horsley, G.-X. Du, M. Pellaton, C. Affolderbach, G. Mileti, and P. Treutlein, *Imaging of relaxation times and microwave field strength in a microfabricated vapor cell*, Physical Review A **88**, 063407 (2013).

[236] J. D. Francesco, F. Gruet, C. Schori, C. Affolderbach, R. Matthey, G. Mileti, Y. Salvadé, Y. Petremand, and N. D. Rooij, *Evaluation of the frequency stability of a VCSEL locked to a micro-fabricated rubidium vapour cell*, in *Semiconductor Lasers and Laser Dynamics IV*, volume 7720, p. 77201T, International Society for Optics and Photonics (SPIE) (2010).

[237] R. Straessle, M. Pellaton, C. Affolderbach, Y. Pétremand, D. Briand, G. Mileti, and N. F. de Rooij, *Microfabricated alkali vapor cell with anti-relaxation wall coating*, Applied Physics Letters **105**, 043502 (2014).

[238] S. Woetzel, E. Kessler, M. Diegel, V. Schultze, and H.-G. Meyer, *Low-temperature anodic bonding using thin films of lithium-niobate-phosphate glass*, Journal of Micromechanics and Microengineering **24**, 095001 (2014).

[239] R. Mhaskar, S. Knappe, and J. Kitching, *A low-power, high-sensitivity micromachined optical magnetometer*, Applied Physics Letters **101**, 241105 (2012).

[240] A. Horsley and P. Treutlein, *Frequency-tunable microwave field detection in an atomic vapor cell*, Applied Physics Letters **108**, 211102 (2016).

[241] L. Weller, *Absolute absorption and dispersion in a thermal Rb vapour at high densities and high magnetic fields*, Ph.D. thesis, Durham University (2013).

[242] F. S. Ponciano-Ojeda, F. D. Logue, and I. G. Hughes, *Absorption spectroscopy and Stokes polarimetry in a ^{87}Rb vapour in the Voigt geometry with a 1.5 T external magnetic field*, Journal of Physics B: Atomic, Molecular and Optical Physics **54**, 015401 (2020).

[243] B. A. Olsen, *Optical pumping and spectroscopy of alkali vapors in magnetic fields*, Ph.D. thesis, Princeton University (2011).

[244] D. J. Whiting, J. Keaveney, C. S. Adams, and I. G. Hughes, *Direct measurement of excited-state dipole matrix elements using electromagnetically induced transparency in the hyperfine Paschen-Back regime*, Physical Review A **93**, 043854 (2016).

[245] C. R. Higgins and I. G. Hughes, *Electromagnetically induced transparency in a V-system with ^{87}Rb vapour in the hyperfine Paschen-Back regime*, Journal of Physics B: Atomic, Molecular and Optical Physics **54**, 165403 (2021).

[246] A. Sargsyan, R. Mirzoyan, and D. Sarkisyan, *Splitting of the electromagnetically induced transparency resonance on ^{85}Rb atoms in strong magnetic fields up to the Paschen-Back regime*, JETP Letters **96**, 303 (2012).

[247] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-photon interactions*, (John Wiley & Sons, Ltd) (1998).

[248] P. M. Anisimov, J. P. Dowling, and B. C. Sanders, *Objectively discerning Autler-Townes splitting from electromagnetically induced transparency*, Physical Review Letters **107**, 163604 (2011).

[249] T. Y. Abi-Salloum, *Electromagnetically induced transparency and Autler-Townes splitting: Two similar but distinct phenomena in two categories of three-level atomic systems*, Physical Review A **81**, 053836 (2010).

[250] X. Lu, X. Miao, J. Bai, L. Pei, M. Wang, Y. Gao, L.-A. Wu, P. Fu, R. Wang, and Z. Zuo, *Transition from Autler-Townes splitting to electromagnetically induced transparency based on the dynamics of decaying dressed states*, Journal of Physics B: Atomic, Molecular and Optical Physics **48**, 055003 (2015).

[251] C. J. Foot, *Atomic physics* (Oxford University Press) (2005).

[252] D. J. Whiting, R. S. Mathew, J. Keaveney, C. S. Adams, and I. G. Hughes, *Four-wave mixing in a non-degenerate four-level diamond configuration in the hyperfine Paschen-Back regime*, Journal of Modern Optics **65**, 713 (2018).

[253] D. Budker and M. Romalis, *Optical magnetometry*, Nature Physics **3**, 227 (2007).

[254] G. Buser, *Storing single photons in broadband vapor cell quantum memories*, Ph.D. thesis, Universität Basel (2021).

[255] JENOPTIK Optical Systems GmbH, Jena, Germany, *Integrated-optical modulators: Technical information and instructions for use* (2019), https://www.jenoptik.com/-/media/websitedocuments/optics/modulators/modulator-manual.pdf.

[256] D. T. Bui, C. T. Nguyen, I. Ledoux-Rak, J. Zyss, and B. Journet, *Instrumentation system for determination and compensation of electro-optic modulator transfer function drift*, Measurement Science and Technology **22**, 125105 (2011).

[257] J. N. Walpole, *Semiconductor amplifiers and lasers with tapered gain regions*, Optical and Quantum Electronics **28**, 623 (1996).

[258] D. C. Kilper and R. S. Tucker, *Chapter 17 – Energy-efficient telecommunications*, in I. P. Kaminow, T. Li, and A. E. Willner (editors), *Optical Fiber Telecommunications*, Optics and Photonics, pp. 747–791 (Academic Press, Boston), sixth edition (2013).

[259] P. Palittapongarnpim, A. MacRae, and A. I. Lvovsky, *Note: A monolithic filter cavity for experiments in quantum optics*, Review of Scientific Instruments **83**, 066101 (2012).

[260] A. Ahlrichs, C. Berkemeier, B. Sprenger, and O. Benson, *A monolithic polarization-independent frequency-filter system for filtering of photon pairs*, Applied Physics Letters **103**, 241110 (2013).

[261] C. M. Natarajan, M. G. Tanner, and R. H. Hadfield, *Superconducting nanowire single-photon detectors: Physics and applications*, Superconductor Science and Technology **25**, 063001 (2012).

[262] G. N. Gol'tsman, O. Okunev, G. Chulkova, A. Lipatov, A. Semenov, K. Smirnov, B. Voronov, A. Dzardanov, C. Williams, and R. Sobolewski, *Picosecond superconducting single-photon optical detector*, Applied Physics Letters **79**, 705 (2001).

[263] J. Yang, A. Kerman, E. Dauler, V. Anant, K. Rosfjord, and K. Berggren, *Modeling the electrical and thermal response of superconducting nanowire single-photon detectors*, IEEE Transactions on Applied Superconductivity **17**, 581 (2007).

[264] I. Esmaeil Zadeh, J. Chang, J. W. N. Los, S. Gyger, A. W. Elshaari, S. Steinhauer, S. N. Dorenbos, and V. Zwiller, *Superconducting nanowire single-photon detectors: A perspective on evolution, state-of-the-art, future developments, and applications*, Applied Physics Letters **118**, 190502 (2021).

[265] M. A. Rosenberry, J. P. Reyes, D. Tupa, and T. J. Gay, *Radiation trapping in rubidium optical pumping at low buffer-gas pressures*, Physical Review A **75**, 023401 (2007).

[266] S. E. Thomas, J. H. D. Munns, K. T. Kaczmarek, C. Qiu, B. Brecht, A. Feizpour, P. M. Ledingham, I. A. Walmsley, J. Nunn, and D. J. Saunders, *High efficiency Raman memory by suppressing radiation trapping*, New Journal of Physics **19**, 063034 (2017).

[267] J. A. Bellisio, P. Davidovits, and P. J. Kindlmann, *Quenching of rubidium resonance radiation by nitrogen and the noble gases*, The Journal of Chemical Physics **48**, 2376 (1968).

[268] M. A. Perez, U. Nguyen, S. Knappe, E. A. Donley, J. Kitching, and A. M. Shkel, *Rubidium vapor cell with integrated Bragg reflectors for compact atomic MEMS*, Sensors and Actuators, A: Physical **154**, 295 (2009).

[269] R. Chutani, V. Maurice, N. Passilly, C. Gorecki, R. Boudot, M. Abdel Hafiz, P. Abbé, S. Galliou, J. Y. Rauch, and E. De Clercq, *Laser light routing in an elongated micromachined vapor cell with diffraction gratings for atomic clock applications*, Scientific Reports **5**, 14001 (2015).

[270] H. Nishino, Y. Yano, M. Hara, M. Toda, M. Kajita, T. Ido, and T. Ono, *Reflection-type vapor cell for micro atomic clocks using local anodic bonding of 45° mirrors*, Optics Letters **46**, 2272 (2021).

[271] C. Tomschitz, *A photoionization scheme to create cold ionic impurities from Rydberg atoms*, Master's thesis, Universität Stuttgart (2018).

[272] M. Yan, E. G. Rickey, and Y. Zhu, *Electromagnetically induced transparency in cold rubidium atoms*, Journal of the Optical Society of America B **18**, 1057 (2001).

[273] P. Walther, M. D. Eisaman, A. André, F. Massou, M. Fleischhauer, A. S. Zibrov, and M. D. Lukin, *Generation of narrow-bandwidth single photons using electromagnetically induced transparency in atomic ensembles*, International Journal of Quantum Information **05**, 51 (2007).

[274] I. Vurgaftman and M. Bashkansky, *Suppressing four-wave mixing in warm-atomic-vapor quantum memory*, Physical Review A **87**, 063836 (2013).

[275] S. Knappe, *3.18 - MEMS atomic clocks*, in Y. B. Gianchandani, O. Tabata, and H. Zappe (editors), *Comprehensive Microsystems*, volume 3, pp. 571–612 (Elsevier, Oxford) (2008).

[276] D. Salart, O. Landry, N. Sangouard, N. Gisin, H. Herrmann, B. Sanguinetti, C. Simon, W. Sohler, R. T. Thew, A. Thomas, and H. Zbinden, *Purification of single-photon entanglement*, Physical Review Letters **104**, 180504 (2010).

[277] F. Gruet, M. Pellaton, C. Affolderbach, T. Bandi, R. Matthey, and G. Mileti, *Compact and frequency stabilized laser heads for rubidium atomic clocks*, in B. Cugny, E. Armandillo, and N. Karafolas (editors), *International Conference on Space Optics - ICSO 2012*, volume 10564, p. 105642Y (SPIE) (2017).

[278] M. D. Rotondaro and G. P. Perram, *Collisional broadening and shift of the rubidium D_1 and D_2 lines by rare gases, H_2, D_2, N_2, CH_4 and CF_4*, Journal of Quantitative Spectroscopy and Radiative Transfer **57**, 497 (1997).

[279] S. Bize, Y. Sortais, M. S. Santos, C. Mandache, A. Clairon, and C. Salomon, *High-accuracy measurement of the ^{87}Rb ground-state hyperfine splitting in an atomic fountain*, Europhysics Letters (EPL) **45**, 558 (1999).

[280] C. Affolderbach, A. Nagel, S. Knappe, C. Jung, D. Wiedenmann, and R. Wynands, *Nonlinear spectroscopy with a vertical-cavity surface-emitting laser (VCSEL)*, Applied Physics B **70**, 407 (2000).

[281] E. Arimondo, *V Coherent population trapping in laser spectroscopy*, in E. Wolf (editor), *Progress in Optics*, volume 35, pp. 257–354 (Elsevier) (1996).

[282] S. Karlen, *Fabrication and characterization of MEMS alkali vapor cells used in chip-scale atomic clocks and other atomic devices*, Ph.D. thesis, Université de Neuchâtel (2017).

[283] J. Vanier, R. Kunski, N. Cyr, J. Y. Savard, and M. Têtu, *On hyperfine frequency shifts caused by buffer gases: Application to the optically pumped passive rubidium frequency standard*, Journal of Applied Physics **53**, 5387 (1982).

[284] J. Vanier and C. Audoin, *The quantum physics of atomic frequency standards* (IOP Publishing Ltd) (1989).

Acknowledgments

Many persons have contributed over the course of my work to the realization of this thesis.

I want to express the deepest gratitude towards Philipp Treutlein who has been extremely supportive of my work. I had the privilege to learn a lot from his expertise. I am thankful for his mentoring over these years, while giving me a lot of freedom. I am also grateful for the many opportunities he rendered possible. My thanks go also to Richard Warburton. I always enjoyed the stimulating discussions we had over the years about our collaborative work. I'm looking forward to finally interfacing the two platforms we have been progressively making more and more compatible. My sincere thanks also go to Mikael Afzelius for agreeing to act as external referee for this work.

I thank Andrew Horsley for his supervision during the early phases of my work. He has been my guide at the beginning of my journey into the Paschen-Back regime. Many thanks go to Janik Wolters, who was of significant help with the basis of the experimental work I presented in this thesis. Even after he moved to Berlin, he was available for discussion about memories and downconversion. I also appreciate that he helped me find a place to stay when I first moved to Basel. It has been a great pleasure and was very helpful to work with Gianni Buser over the last few years as a 'lab buddy'. I always appreciated his company in the lab, his choice of music, but also the occasional board-game night and the fact that he started supplying me with home-pickled vegetables. I also thank him for the meticulous work he did while proofreading this thesis. It has been a great pleasure supervising Björn Cotting. He stayed motivated even when Gianni and I gave him impossible tasks. His work has proven very fruitful for our following experiments. The rest of the Treutlein group has my appreciation for the constructive discussions but also for the interesting conversation topics during our coffee breaks. My thanks go to the current team: Madhavakkannan Saravanan, Manel Bosch Aguilera, Maryse Ernzer, Gian-Luca Schmid, Paolo Colciaghi, Tilman Zibold, Yifan Li, and Yongqi Shi; but also to former members: James Ngai, Matteo Fadel, Thomas Karg, Boris Décamps, Baptiste Gouraud, Andreas Jöckel, and Aline Vochezer.

Furthermore, I wish to thank Gaetano Mileti and his group in Neuchâtel for providing me with the MEMS vapor cell I used for the HPB memory, without which, a big portion of the work presented here would not have been possible. My thanks go especially to Florian Gruet for the time he spent helping me to characterize the cell and reverse engineer its content.

I benefited a lot from the collaborative work and the exchange of expertise (but also equipment) with Nicolas Sanguouard's, Richard Warburton's and Oliver Benson's group. Specifically, I want to thank Pavel Sekatski for our discussions about $g^{(2)}$'s, Liang Zhai for providing me with SPADs, Alisa Javadi and Natasha Tomm for the help in the jitter determination and Josh Zuber for his help with the single-photon spectroscopy.

I wish to thank Sascha Martin and the whole team of the mechanical workshop for the fabrication of custom-made components, even when only on a short notice. My CAD-skills have significantly improved due to your feedback on my designs. I am grateful to Dominik Sifrig for repairing the water cooling of the power supply of the electromagnet, allowing me to perform the last measurements described in this thesis. I thank Michael Steinacher and the whole team from the electronics workshop for resuscitating the power supply of the electromagnet (more than once), designing the current switches for the SOAs, and for helping with more mundane things like supplying us with connectors or advice on how to solder properly.

And last, but not least, thanks mom and dad, especially for your patience.

(96257) Roberto = 1995 JE
Discovery: 1995-05-03 / S. Mottola / La Silla / 809

Roberto Mottola (b. 1990) is an Italian-German physicist expert in the field of quantum physics. He has investigated the use of room-temperature single-photon sources and atomic quantum memories for the application to broadband quantum networks. Name proposed by his father, the discoverer, in occasion of the successful completion of his PhD in March 2023.

Excerpt from: WG Small Bodies Nomenclature Bulletin **3**, #4 (20 March 2023), ISSN 2789-2603, published on behalf of the International Astronomical Union.