

# Storing Single Photons in Broadband Vapor Cell Quantum Memories

Inauguraldissertation

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

von

**Gianni Carlo Buser**

2023

Originaldokument gespeichert auf dem Dokumentenserver der Universität Basel  
<http://edoc.unibas.ch>



This work is dedicated to the public domain via CC0  
The complete text of the CC0 deed may be reviewed at  
<https://creativecommons.org/publicdomain/zero/1.0/>

Genehmigt von der Philosophisch-Naturwissenschaftlichen Fakultät  
auf Antrag von

Erstbetreuer: Prof. Dr. Philipp Treutlein

Zweitbetreuer: Prof. Dr. Richard Warburton

externer Experte: Prof. Dr. Ofer Firstenberg

Basel, den 14. Dezember 2021

Prof. Dr. Marcel Mayor  
Dekan

# Abstract

Single photons are an essential resource for realizing quantum technologies. Together with compatible quantum memories granting control over when a photon arrives, they form a foundational component both of quantum communication and quantum information processing. Quality solid-state single photon sources deliver on the high bandwidths and rates required for scalable quantum technology, but require memories that match these operational parameters. In this thesis, I report on quantum memories based on electromagnetically induced transparency and built in warm rubidium vapor, with such fast and high bandwidth interfaces in mind. I also present work on a heralded single photon source based on parametric downconversion in an optical cavity, operated in a bandwidth regime of a few 100s of megahertz. The systems are characterized on their own and together in a functional interface. As the photon generation process is spontaneous, the memory is implemented as a fully reactive device, capable of storing and retrieving photons in response to an asynchronous external trigger.

The combined system is used to demonstrate the storage and retrieval of single photons in and from the quantum memory. Using polarization selection rules in the Zeeman substructure of the atoms, the read-out noise of the memory is considerably reduced from what is common in ground-state storage schemes in warm vapor. Critically, the quantum signature in the photon number statistics of the retrieved photons is successfully maintained, proving that the emission from the memory is dominated by single photons. We observe a retrieved single-photon state accuracy of  $g_{c,\text{ret}}^{(2)} = 0.177(23)$  for short storage times, which remains  $g_{c,\text{ret}}^{(2)} < 0.5$  throughout the memory lifetime of 680(50) ns. The end-to-end efficiency of the memory interfaced with the photon source is  $\eta_{e2e} = 1.1(2)\%$ , which will be further improved in the future by optimizing the operating regime. With its operation bandwidth of 370 MHz, our system opens up new possibilities for single-photon synchronization and local quantum networking experiments at high repetition rates.



# Acknowledgments

Many people lent their hands and minds to aid this work. I am grateful to them. I thank Philipp Treutlein, from whom I have had the pleasure of learning since the third semester of undergrad. Had me hooked early on, I guess. Within the context of working in his group for some 6 years, I thank him for great freedoms and general support, as well as regular, more direct inspiration. Whether the topic was experiments, theory, justice, anarchy, or teaching, I have never had a fruitless discussion with him, and I treasure his trust. I also thank Richard Warburton. His wit and poignancy always refocused our collaborative plans of interfacing dots and atoms on the next step towards the hope of measuring anything at all worthwhile, and while we did not quite get there yet our road is more open than ever. I thank Janik Wolters, who features a few more times in this text, for some of the most productive, continuous laboratory arguments imaginable. I learned a lot from him, and the process was fun. I thank Roberto Mottola for many things, including some proofing, and our work has become sufficiently connected that to be safe I better just thank him for his company. I thank Björn Cotting for his labors on “source v2”. It was a delight guiding him away from the mistakes Roberto and I made and into some new ones.

I am appreciative towards my many colleagues in the Treutlein group throughout the years for their reliable feedback cycles, and their cakes. I thank Thomas Karg, Matteo Fadel, Tilman Zibold, Maryse Ernzer, Paolo Colciaghi, Yifan Li, James Ngai, Gian-Luca Schmid, Yongqi Shi, and Manel Bosch Aguilera. From a past generation of Treutleins, Lucas Béguin, Baptiste Gouraud, Andreas Jöckel, Aline Vochezer, Boris Décamps, and Roman Schmied also have my gratitude for their contributions to my initial integration. I am looking forward to working more with Madhavakkannan Saravanan.

My work has benefited from many collaborative interactions in the exchange of materials and expertise with diverse people from the groups of Richard Warburton, Nicolas Sanguouard, and Oliver Benson. In particular I want thank Liang Zhai for aid with photon detectors generally, as well as Natasha Tomm and Alisa Javadi for assistance in determining their jitter. I also thank Pavel Sekatski for elucidating his models of coherence functions for a considerable time until we slowly converged on a

---

shared reality. Moreover, I want to thank Michael Steinacher, Andreas Tonin, and the whole team of the electronics lab for their regular help with a connector or crimp, and in particular for the design and construction of the SOA driver. Sasha Martin and the team of the mechanics shop also have my thanks for many a part made and threading turned.

I thank Kevin Saladin and Carlo Seppi for their diverse inputs on the introduction. As a stand in for many more personal words of thanks to my friends I single out Mischa Frech to thank him for his cheeky invention of the Frech constant. I assure you, dear reader, that I have been scrupulous in eliminating it from the final manuscript. Finally, there are my parents, who supported me and contributed to so many aspects of my well-being during these amazing years. Thanks mom and thanks dad.

One more note: I thank the people in this world who believe in free and accessible information, and choose to act accordingly. As symbolic tribute, I dedicate my thesis to the public domain. Intellectual property is theft from the heritage of humanity. No borders for information. No borders for people.

# Contents

<b>Abstract</b>	<b>i</b>
<b>Acknowledgments</b>	<b>iii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 The Landscape of Classical Memories . . . . .	2
1.2 Definitions and Distinctions: Enter the Quantum . . . . .	5
1.2.1 Storing Quantum Information . . . . .	7
1.2.2 Figures of Merit . . . . .	9
<b>2 Review of Memories for Light and Single Photon Sources</b>	<b>17</b>
2.1 Review of Memories for Light . . . . .	18
2.1.1 Delay Lines . . . . .	18
2.1.2 Optimal Cloning . . . . .	20
2.1.3 Approaches to Optical Depth . . . . .	20
2.1.4 Lambda Memories . . . . .	23
2.1.5 Ladder Memories . . . . .	41
2.1.6 Beyond Three-Level Systems and Nuclear Coherence . . . . .	42
2.2 Review of Single Photon Sources . . . . .	43
2.2.1 Figures of Merit . . . . .	45
2.2.2 Atomic Sources and Read-Only Memories . . . . .	50
2.2.3 Semiconductor Quantum Dots . . . . .	51
2.2.4 Spontaneous Parametric Down-Conversion . . . . .	53
2.3 Overview of Applications . . . . .	54
2.3.1 Quantum Communication: Infrastructure Projects . . . . .	55
2.3.2 Quantum Communication: Cryptography . . . . .	57
2.3.3 Quantum Computing Optically . . . . .	60
2.3.4 Quantum Computing on a Quantum Network . . . . .	62
<b>3 Hyperfine Memory</b>	<b>65</b>
3.1 Scheme and Implementation . . . . .	66

3.2	Pulse Generation . . . . .	69
3.2.1	Electo-Optic Modulators . . . . .	70
3.2.2	Arbitrary Waveform Generator . . . . .	73
3.3	Optical Setup . . . . .	75
3.3.1	Signal Pulse Calibration . . . . .	76
3.3.2	Control Pulse Amplification, Optical Pumping . . . . .	79
3.3.3	Vapor Cell and Memory . . . . .	80
3.3.4	Filtration . . . . .	82
3.3.5	Auxiliary Setup Elements . . . . .	85
3.4	Results . . . . .	86
3.5	Performance Optimization and Noise Analysis . . . . .	91
3.6	Towards Storing Single Photons . . . . .	94
<b>4</b>	<b>SPDC Photon Pair Source</b>	<b>97</b>
4.1	Source Design and Expected Performance . . . . .	98
4.2	Setup and Operation . . . . .	104
4.3	Output Characterization . . . . .	107
4.4	Limitations and Double-Resonant Version . . . . .	112
4.5	Alterations for Interfacing . . . . .	114
4.6	Tuning to the Rb D2 Line . . . . .	116
<b>5</b>	<b>Zeeman Memory</b>	<b>119</b>
5.1	Scheme . . . . .	119
5.1.1	Line Shapes and Excited State Interference . . . . .	121
5.2	Setup . . . . .	123
5.2.1	Requirements of an Externally Triggered Memory . . . . .	124
5.2.2	Control Pulse Amplification . . . . .	132
5.2.3	Switching of the Optical Pumping Mode . . . . .	136
5.2.4	Vapor Cell and Memory . . . . .	139
5.2.5	Filtration . . . . .	141
5.3	Single Photon Storage . . . . .	144
5.4	Further Characterization . . . . .	149
5.4.1	Determining the Atomic Temperature . . . . .	150
5.4.2	State Preparation . . . . .	151
5.4.3	Noise Spectra . . . . .	154
<b>6</b>	<b>Outlook</b>	<b>159</b>
<b>7</b>	<b>Appendices</b>	<b>163</b>
7.1	Full Level Diagram for Hyperfine Memory . . . . .	163



*Contents*

---

7.2	HBT Detection Jitter Measurement . . . . .	165
7.3	Absolute Frequency Calibration by Rb Spectroscopy . . . . .	165
7.4	Hyperfine Memory Noise Spectrum . . . . .	167
	<b>List of Figures</b>	<b>169</b>
	<b>List of Tables</b>	<b>170</b>
	<b>Bibliography</b>	<b>171</b>



# Chapter 1

## Introduction

*Between my finger and my thumb  
The squat pen rests; snug as a gun.*

Seamus Heaney

---

Memories, at their core, are best thought of as enabling tools. They are building blocks of network architecture, functioning as buffers and repeaters, and essential components of computers in the form of working memories for fast access by a processor. More fundamentally, and free from the context of computers, they are a means of bestowing permanence to the ephemeral acting as non-volatile information storage. This is true of physically classical memories, where types are so diverse and applications so widespread that the technological aspect often becomes invisible – consider that a piece of paper is a kind of memory – as well as of quantum memories. I will discuss definitions and the limits of these comparisons in more detail later in this chapter; until then let a memory be a tool that solves the problem of information being lost to its user.<sup>1</sup> Applying this framework to the context of quantum optics and single photons specifically, I believe the ideal future aspiring memories should strive for is to become simple, cheap, and ubiquitous additions to experiments in this realm. Preferably, to be taken for granted as straightforward problem-solving devices, neither particularly complicated nor practically inconvenient, perhaps about as demanding as home-built diode lasers today.

While not precluding such uses, this picture of quantum memories as tools is not to imply a focus on near-term technological applications in the wider world. Research and tools go hand in hand, as so often it is not a lack of ideas but the missing means to implement them that bottlenecks experimental results. Advances on the

---

<sup>1</sup>I'm treading carefully here due to the conservation of quantum information.


level of tools open the door to everything else. Theodor Hänsch, whose early work includes developing narrow linewidth lasers [1], illustrates this point beautifully when he continues characterizing his later work on frequency combs as revolving around making tools in his Nobel lecture [2]. This backdrop of functionality is a recurring theme in this introduction, and motivates both the discussion of figures of merit in the context of the applications of quantum memories, as well as the analogies and differentiation between them and their classical counterparts.

My focus will then narrow down to the main topics of interest in this thesis, solid-state single photon sources and quantum memories for single photons. In chapter 2, I will review the literature on memories and single photon sources to set the stage for the following experimental chapters. In chapter 3, I will first describe initial experiments with a simple memory, characterized with laser pulses, implemented in the hyperfine split ground states of warm rubidium vapor. This experiment defines the requirements of a compatible photon source, which is then designed to meet them. I describe this process and present a characterization of photons generated by parametric downconversion in chapter 4. Finally, in chapter 5, I detail experiments interfacing this photon source with a refined memory scheme novel to this application. As a key result, I report on the storage and retrieval of single photons in a warm vapor quantum memory. The results constituting this thesis are also published in

1. 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).
2. 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  $^{87}\text{Rb}$  D<sub>1</sub> line*, Optics Express **28**, 3159 (2020).
3. 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).

## 1.1 The Landscape of Classical Memories

Since the beginning of the information age the quest for ever better physical hardware implementations of processors, memories, and interfaces has been ongoing. For processors, the overwhelming success of the silicon transistor has left little room for viable competition in the preceding decades. This does not necessarily mean that the last word has been spoken on that front. Success has hinged on Moore’s Law [3], an uncannily successful prediction of the exponential growth of the number of components that can



**Figure 1.1:** The memory hierarchy in a typical computer. As the gradient lightens upwards the distance to the CPU increases, corresponding to longer access times as well as more readily available storage space, typically limited by cost.

be fit on a chip, which is predicated on scaling and miniaturization. Manufacturing has already gently touched upon the limit of single-atom devices [4], and modern FinFETs in commercially produced devices are not much larger, nor is that the only fundamental limit to consider [5]. It stands to reason that these racing improvements will stop and room may emerge for technologies that offer more favorable fundamental limits or other advantages, such as oft written off all-optical processors [6], or nascent carbon-nanotube (CNT) devices [7, 8], assuming they can achieve competitive scaling once the gap stops widening.

Dissimilarly for memories, even in the narrower, computer-oriented sense, this search has resulted in a menagerie of specialized implementations. The needs of a long-term data archive are sufficiently different from those of a supercomputer’s CPU that, while both require some kind of memory, a panacea that perfectly satisfies the quality criteria and budgets of both simultaneously seems fundamentally implausible. In practice, the distinction usually boils down to a performance vs cost trade-off – speed is expensive. Additionally, the permanence of the storage solution also plays an important role, but is considered separately as historically it generally anti-correlates with performance in terms of speed.

Even within a regular home or office computer this trade-off yields a memory hierarchy [9], drawn schematically in figure 1.1. A small amount of the fastest and most expensive memory, generally static random-access memory (SRAM) storing bits in flip-flops, is used for the processor’s cache. The scales at play here are on the order of hundreds of picoseconds access times, tens of megabytes of storage space, and 6–8 transistors required per bit [10]. One level down the hierarchy we find dynamic random-access memory (DRAM), now pairing one transistor with one capacitor per bit. The DRAM is still fast and holds a much larger amount of data close at hand to feed up the chain to the processor. These types of memory are volatile, requiring power to maintain their state, although not sufficiently volatile to avoid becoming security hazards [11]. These are two more features to consider when discussing the attributes of memories: upkeep costs and physical security. Random-access memory development is open to new physical implementations, seeking higher memory densities,

lower power consumption per read/write operation, and non-volatility even at RAM speeds. Interesting modern approaches include anti-ferromagnetic magnetoelectric [12] (AF-MERAM) and magnetoresistive [13] (AF-MRAM) random access memories, the exploitation of anti-ferromagnetic order constituting the latest twist on older magnetoresistive memory ideas [14]. Next down the memory hierarchy come the local storage disks, where there is ongoing competition between faster solid-state flash memory implementations and higher capacity per unit cost magnetic hard disks. Optical discs also belong in this category, where yet another physical implementation of storage underlies. Categorically important is that the information is still stored locally, whereas the distinction between a drive inside the computer case and a disc on the bookshelf is insignificant. Finally, on the last rung of the hierarchical ladder, there are the needs of servers, archives, and cloud services to consider. For these information agglomerations in remote locations, serving many people, the required capacities are growing huge: the scale now approaches zettabytes [15]. Arguably, no current technology is truly suitable to cover this need, and research into new alternatives, for example storage in fused silica [16], is ongoing. So as not to conflate this hierarchy with a general quality measure, and to return the discussion to tools with a purpose, consider that the volatility of SRAM memories would be burdensome, the speed would be useless, and the cost would be stupefying for the *goal* of archiving data.

Finer resolutions could be applied to this overview. There are interesting historical memory implementations, many more in development, as well as numerous specific measures to use as proxies for theoretical best performances when comparing different physical implementations. From among these proxies the physical density of bits is often cited, as it has implications for access times once the limiting factor becomes how far the information carriers must travel, as well as for the device size required for mass storage. Some more implementations in development include spintronic research into using magnetic domain walls in nanowires to store information [17], carbon-based devices that use the relative positions (near, far) of CNT pairs and the associated resistance (low, high) across them as bits [18], and a nearly endless list of devices with ever more esoteric advantages. For those enamored with the natural, an excursion into the realms of DNA memory [19] might also be worthwhile. Fortunately, such comprehensive granularity is not required to conclude that the field of classical memories is diverse, specialized, and highly active.

This final point allows us to predict something concrete about the future of quantum memories. As tools, classical memories have gotten ever more adapted to specific, narrow purposes. “Going quantum” does not remove trade-offs from the equation. Indeed, we are currently in the teething phase, where there are many more trade-offs in basic functionality to consider. Even when these are overcome, I believe that the maturity of quantum memory technology will be signaled by a wide variety of

implementations, and not by a monolith. This motivates a detailed look at a broad spectrum of figures of merit in section 1.2. First, however, let us get more specific with definitions and the distinctions to be made between classical and quantum memories.

## 1.2 Definitions and Distinctions: Enter the Quantum

Conceptually, although not in name, the bit appeared in thermodynamic literature [20] some twenty years before Shannon's seminal work in 1948 [21] considered the founding of information theory [22]. The link was made explicit by his naming the measure of information contained in a random variable  $X$  its entropy. For a discrete probability distribution  $p_1, \dots, p_n$  the Shannon entropy in bits is defined as,

$$H(X) = - \sum_x p_x \log_2 p_x. \quad (1.1)$$

Computers have habituated us to quantifying data in bits, and their information content is reasonably tangible. The bit's two states can be thought of as true and false, encoding the answers to logical propositions. Human-readable representations of information, like text, generally use an *alphabet* with more than two symbols, or may even come in analog form, like speech, which can be thought of as the large alphabet limit. Indeed the choice of the bit is not uniquely suited up to a change of logarithmic basis, and the sensible choice is the number of symbols in the alphabet of the representation of the information. The most natural choice of alphabet for physical encoding varies with the degrees of freedom of the system used to represent the information, bit for two, trit for three, quadrit for four, and so on. Physical signals can also be modeled as continuous variables. This requires information measures with slightly different definitions but analogous purposes and the same names [23]. When adjusting for the implementation in this way, the entropy directly accounts for the physical resources required for an optimal encoding.<sup>2</sup> Usually longer alphabets imply more complicated devices, up to human-level complexity for logographic scripts with thousands of symbols, so while classical [25] and quantum mechanical [26] implementations of ternary logic are physically studied, larger ones are rarer and the common workhorse is two dimensional.

Analogously a quantum system with 2 orthogonal states linked by a uniquely addressable transition can be considered the physical implementation of a qubit. A

---

<sup>2</sup>Optimal meaning no information is lost and no resources are wasted. Perfectly compressed. The issue arises as using a biased random variable as an alphabet reduces the information content per symbol. Consider, for example, that excluding loanwords the letter q is always followed by u in English writing, causing encoding redundancy. Shannon himself investigated the redundancy of printed English in some detail [24] – naturally results vary with text-type, and exact values are impossible to state generally, but 54% appears to be a reasonable lower bound.

## 1.2. Definitions and Distinctions: Enter the Quantum

---

pure state of this system can always be written as,

$$|\Psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad (1.2)$$

where  $|0\rangle$  and  $|1\rangle$  are our free choice of orthonormal basis states, and  $\alpha$  and  $\beta$  are the corresponding probability amplitudes. When describing an ensemble in a mixture of states  $|\Psi_i\rangle$  with probabilities  $p_i$ , or just for the sake of generality even when the state is pure, we can describe it with the density matrix,

$$\rho = \sum_i p_i |\Psi_i\rangle \langle \Psi_i|. \quad (1.3)$$

How does this object compare to a classical bit? At a glance, it appears to contain an arbitrarily large amount of information. After all, the coefficients in equation 1.2 are complex numbers and anything could have been encoded into the digits of those coefficients. It turns out however, that the information accessible through measurement is limited by Holevo's bound [27] to a maximum of one bit, although it is possible to access less. Indeed, this possibility is to be expected, as there is no measurement that reliably distinguishes non-orthogonal quantum states [28]. There are two further aspects of the qubit that distinguish it from a classical bit right away. First, superposition states are allowed, which also implies that there is the meaningful possibility of measuring in multiple bases. Second, there is the effect that a measurement changes the qubit state, rather than just changing our knowledge about the state. The property that will allow us to do more with quantum states than is possible classically is called entanglement and has to do with the correlations supported between systems, but it should also be stated right away that this property alone does not exclude efficient classical simulation of the system [29].

Traditionally, there is a class of states treated as “free” in quantum information theory [30]. These states, called the separable states, take the form

$$\rho = \sum_k p_k \rho_1^k \otimes \dots \otimes \rho_N^k \quad (1.4)$$

for  $N$  subsystems or parties, with states  $\rho_i$  acting on the Hilbert space  $H_i$  ( $i = 1 \dots N$ ), and the full density matrix acting on  $H = H_1 \otimes \dots \otimes H_n$ . Even if the parties are spatially separated from one another, these states can be created by protocols called local operations and classical communication, or LOCC. Classical communication is self explanatory practically, and even its theoretical abstraction of the exchange of bits comes naturally today. The local operations refer to arbitrary measurements, rotations, and interactions done by one party on their own subsystem. In practice, it may of course turn out that these local operations are difficult to realize, however, nothing



therein will defy a classical worldview. This observation corresponds to two possible definitions of entanglement. The first and most famous definition is that entangled states are not mixtures of product states – they are not separable. Irritatingly, this definition is negative and thus says nothing about what entangled states actually are. An alternative is to define an entangled state as a useful resource for completing a non-classical task [31]. This definition is more modern and somewhat obscurer than the former, as it was only shown in 2006 that truly any quantum correlations give an information-processing advantage over classical correlations [32]. Nevertheless, when viewed as a resource, the question of how much entanglement there is arises naturally, and our positive definition offers no more immediate guidance on that matter than its negative counterpart.

The quantum equivalent to the Shannon entropy is called the von Neuman entropy, defined for a state  $\rho$  as

$$S(\rho) = -\text{tr}(\rho \log_2 \rho). \quad (1.5)$$

Taking for granted that the non-classical correlations within this state are what interest us, how are these quantified? There are many options. Considering the case of a pure state with only two subsystems of interest  $A$  and  $B$ , i.e. a bipartite pure state, we can use the mutual information,

$$I(A; B) = S(\rho_A) + S(\rho_B) - S(\rho_{AB}). \quad (1.6)$$

Here  $\rho_A$  is constructed from the full state  $\rho_{AB}$  via a partial trace over  $B$  and vice versa. It can be shown via a Schmidt decomposition of  $\rho$  that this reduces to  $I(A; B) = 2S(\rho_A)$  [28], so, in fact, the entropy can be used directly as an entanglement measure for bipartite pure states. For mixed states of small systems, there is also a sufficient separability criterion involving the partial transpose of  $\rho$  [33]. In general, however, the question of whether a large arbitrary density matrix describes an entangled state is NP-hard [34], and so tailored inequalities excluding all separable states and some amount of entangled states called *witnesses* [35] are usually used to detect entanglement for large systems in practice. That said, for theoretical purposes, any positive real function of a state which is zero if the state is separable and does not increase under LOCC may serve as a measure of entanglement. For a survey of these measures and their applications see [36].

### 1.2.1 Storing Quantum Information

We have identified entanglement as the secret sauce in qubits that makes performing non-classical tasks possible, and we will revisit the nature and value of these tasks in section 2.3. The potential presence of such a property already hints at the fact that

unconventional means of storage will be required to preserve quantum information. The restrictions begin with a famous no-go theorem of quantum physics: the no-cloning theorem. It states that an unknown quantum state cannot be perfectly duplicated [37]. In particular, there exists no unitary operator that can perform the transformation  $|\Psi\rangle|0\rangle \mapsto |\Psi\rangle|\Psi\rangle$ . This rules out many classical techniques and applications. The conventional notion of a backup is impossible, and we cannot seek safety in numbers when the risk of errors is high. Even without invoking the specter of some catastrophic failure causing data loss, a classical computer routinely protects bits through error correction. The simplest picture<sup>3</sup> of this is as follows: for a logical bit of information  $b \in \{0, 1\}$ , instead of storing  $b$  in only one physical bit, a computer stores three copies,  $bbb$ , in three physical bits. Then, when reading out the bit, the logical bit is taken to be the most common value in  $bbb$ . This protects the bit with certainty from one bit-flip error. This simple scheme is not possible for qubits, and although alternative means of correcting errors and achieving fault tolerance are an active and promising theoretical field they have proven to be experimentally challenging [38].

This realization must temper our ambitions. We cannot hope for a quantum copy-and-paste, the best we can do is a quantum cut-and-paste. This may raise the question of whether a quantum memory is useful at all. Fortunately it is simple to quell doubts of this kind, but we must first get more concrete about the physical implementations of qubits and their limitations. To be viable for quantum computation, the physical implementation of a qubit should fulfill a list of minimal criteria, 5 general ones plus 2 for the viability of networking: the DiVincenzo criteria [39]. For most applications of quantum communication, such as quantum cryptography, the list is far shorter – it must be possible to accurately transmit qubits between distant locations. In other words, qubits that can move around, so called flying qubits, are required. While using electrons [40] or even phonons [41, 42] for this purpose is by no means unthinkable, photons are so uniquely well suited to the task due to their speed, limited interactions with the environment, and with one another that the discussion of using something else is generally motivated by its value as a stationary qubit implementation. Photons are the natural choice for transmission [43], and also have the potential to be the medium for computation via a variety of schemes [44, 45, 46, 47]. Reflecting, once again, on the classical computing and communication paradigms that have asserted themselves, it also seems overwhelmingly likely that even if a different physical system is used for computation, light will nevertheless be relied on for communication. Quantum SWAP gates between photons and different physical qubit implementations already exist [48, 49]. Moreover, telecom photons traveling through fibers set a high bar for the quality of novel transmission channels for another kind of flying qubit, and conveniently

---

<sup>3</sup>In practice, error correction algorithms are significantly more sophisticated than this.

the large-scale infrastructure of fiber communication often already exists.

This leads us directly to the needs addressed by quantum memories. Qubits that constantly fly around at the speed of light are great for transmission, but what happens when they arrive? Chances are they will need to be stopped. This is an immediate use scenario of a quantum memory, a temporary cache for quantum information that would otherwise be lost, either due to the qubits literally flying away or due to fast decoherence in some other computational qubit. Such a cache is necessary not just for transmissions, but also to synchronize local probabilistic processes, including some photon sources and quantum gates [50]. Further, as no-cloning forbids amplification of qubit signals in the regular sense, quantum repeater protocols are required to transmit information over lossy channels, loss being inherent to transmission over long distances. These, in turn, typically rely on quantum memories [51, 52]. All these things are the building blocks that constitute the more intrinsically enticing and flashy vision of a quantum internet<sup>4</sup> [53, 54].

### 1.2.2 Figures of Merit

Before we focus in on implementations of quantum memories for photons we must establish measures by which to evaluate the quality of memories. Most of these figures of merit are agnostic towards the physical implementation, or at least have analogues applying to all kinds of memories, including classical ones. These measures will allow us to compare memories, and establish thresholds that must be surpassed for a memory to improve upon the situation without a memory. Some vary in their importance with the target application, and some are specific to memories for photons. As this is an area where a common language is vital, I have stuck to the definitions laid out in [55] where possible, while also discussing a handful of common deviations.

#### Fidelity

The fidelity of a memory measures how closely what is read out matches what was read in. For quantum memories it is given by the state overlap between the input and output states,

$$\mathcal{F} = |\langle \Psi_{\text{in}} | \Psi_{\text{out}} \rangle|. \quad (1.7)$$

For mixed states it takes the form [56],

$$\mathcal{F} = \text{tr} \left( \sqrt{\sqrt{\rho_{\text{in}}} \rho_{\text{out}} \sqrt{\rho_{\text{in}}}} \right). \quad (1.8)$$

---

<sup>4</sup>For the reader seeking concrete information on the *why* of it all I refer to section 2.3 for a brief overview of the applications of quantum information, ceding that practical results may not always be why people do things.

Sensibly, the fidelity is symmetric. In the context of unambiguous state discrimination the fidelity is usually defined as the square of these equations [55]. Note that the question of whether the input and output of a memory are the same applies to truly any kind of storage, from the notoriously low fidelity memory facilitated by neurons to the stone tablet that conceivably retains its fidelity until language itself changes.<sup>5</sup> When there is a chance that no information is retrieved from the memory at all, the fidelity can be conditioned on the success of the read-out to distinguish the chance of a complete failure from the quality degradation induced by successful memory operation. In this case, it is appropriate to refer to the conditional fidelity. These figures are related by the efficiency,  $\mathcal{F} = \eta \mathcal{F}_{\text{cond}}$ . Fidelity is also a useful measure for evaluating the performance of conversions and unitary gate operations represented by an operator  $U$  generally. In such cases it is appropriate to define  $\mathcal{F} = |\langle \Psi_{\text{target}} | \Psi_{\text{out}} \rangle|$ , where  $|\Psi_{\text{target}}\rangle = U |\Psi_{\text{in}}\rangle$  is the target state of the operation. The usual understanding of a memory would then require  $U = \mathbb{1}$ , although it is conceivable that e.g. a rotation or other simple transformation of the output is desirable, or at least easily reversed given that it is known about. In such scenarios, stating a memory fidelity with reference to a  $|\Psi_{\text{target}}\rangle$  is completely appropriate.

### Efficiency

The efficiency of a memory is a useful figure of merit when there is a chance of failure in the retrieval. It facilitates making the distinction between recalling something imperfectly (a fidelity problem) and not recalling anything at all. Particularly for the storage of single photons this is a valuable measure as it is always calculable, even from a simple arrival time histogram. In that case it takes the form,

$$\eta_{\text{mem}} = \frac{N_{\text{out}} - N_{\text{noise}}}{\eta_{\text{det}} N_{\text{in}}}, \quad (1.9)$$

where  $N$  represents the number of photons counted. Often photon memories emit some light upon a read-out attempt even when they are empty. To correct for this effect, a memory characterization without input can be performed, after which the produced number of noise photons  $N_{\text{noise}}$  therein can be subtracted from the counted output during normal operation to yield the real number of retrieved photons. As a measure to characterize only the memory, it is appropriate to correct experimentally

---

<sup>5</sup>There are some famous examples of lost languages and alphabets being rediscovered from deciphering the contents of these kinds of memories, from Egyptian scripts by means of the Rosetta stone, to cuneiform found on Mesopotamian clay tablets, as well as a few memories that seem to have truly outlived their encoding, like the fragmented Linear A tablets that once served the early Minoan civilization. These are of course also some of the most extreme examples of memories with long retention times (see below).

measured values for the efficiencies of the detector, represented by the labeled  $\eta_{\text{det}}$  in the denominator. Moreover, the source may not be perfect either, and the number of actually input photons may differ from the number of attempts by a source efficiency, i.e.  $N_{\text{in}} = \eta_{\text{source}}N_{\text{attempt}}$ .

In theoretical descriptions of memories the success probabilities of read-in and read-out are typically treated separately, therefore the definition of equation 1.9 also takes the name total efficiency. As the efficiency inevitably is a function of the storage time, it is often extrapolated to the limit of zero storage time with the time dependence reported separately (see lifetime). When inefficiencies are due to losses of a known technical nature, these must be accounted for when comparing to theory so as to accurately reflect the physical interaction. The value corrected for such losses is then called the *internal* efficiency – conceptually it is the efficiency of the storage process itself. To contrast, or for emphasis that no such correction was made, our definition can then be referred to as the *external* efficiency or the *end-to-end* efficiency. Without qualifier I will generally use the word only to refer to the latter. Corrections of this kind have the potential to be valuable comparative aids between implementations. That said, care must be taken to distinguish between sources of losses – losses describable as “technical” may be due to filtering of the output for noise, or cavity enhancement of the memory interaction, and therefore are not necessarily avoidable through better coating of optics or other technical solutions. In any real use case, the only interesting efficiency is the external one.

### Lifetime or Retention Time

Memory lifetime measures how long stored information remains retrievable. This measure is an admission that no memory can defeat the second law of thermodynamics; eventually all things must end. Save for the existential kind, there is no need for melancholy at this realization. If a memory is a tool the real question is whether it is useful, and here the relevant comparison is not eternity but how long the information would persist without the memory, or how long it must persist for an application to function. For quantum memories this is usually equated to the coherence time of the storage state  $\tau_c$ , but in practice exact definitions vary and are often adapted to the process by which the information is lost, of which there are many [57]. Practical measures, like the time it takes for another figure of merit most interesting to the conceived application to degrade, are usually the values actually determined experimentally. This could mean the time it takes for the fidelity or the efficiency to fall below a fractional threshold of its starting value (commonly  $1/e$ ). Alternatively how long after storing the memory still provides added value in its application can be specified, e.g. how long an entanglement measure between distant parties is improved by a memory in a quantum

repeater protocol [58]. This application minded definition may seem a little strange, but memory lifetime is always a highly contextual figure of merit – to decide whether a lifetime is good it needs to be compared either to the speed of local operations if used locally, or to the network latency, that is the communication time between nodes, when used in the context of long-distance communication. A measure which is more comparable between different photon storage implementations is the fractional delay. This is the ratio of the achieved storage time and the temporal width of the stored signal.

### **Compatibility: Wavelength and Bandwidth**

It may seem like a trivial observation, but memories have to be compatible with what is to be stored in them. This consideration is relevant on the physical level, the correct voltages need to be applied to flash memories, the correct magnetic fields to hard disks and so on. For photon memories the parameters of interest are the wavelength and bandwidth of the photons to be stored, which must be matched by the memory's storage transition. First and foremost, the matching of these parameters is important, as otherwise the memory will not work at all. Simultaneously, high bandwidth is intrinsically desirable as it determines the operation speed, or more concretely the maximum possible repetition rates. In synchronization applications specifically, where memories are used to enhance the coincidence rates of probabilistic photon sources or gate operations (see section 2.3.3), it is the product of the memory's lifetime and acceptance bandwidth, the time-bandwidth product, scaled by the memory efficiency, that determines the maximal rate enhancement factor [59]. The time-bandwidth product is also a bound on the number of logical operations a quantum processor using the memory could maximally perform [60]. For the application of long-distance quantum communication, fiber losses must be considered as well, placing the optimal wavelengths in the telecom-bands (1550 nm for minimal losses, around 1300 nm for minimal dispersion). If need be, the latter point can also be addressed with frequency conversion, which is rapidly becoming a well established resource [61] with external conversion efficiencies of 57% reported in state-of-the-art experiments [62].

### **Mode Capacity**

In the storage of light, mode capacity can refer to several things. The modes can be literal, for example multiple spatial modes can be stored at once in atomic ensembles in what is called a quantum hologram [63]. This topic involves a great deal of interesting atomic physics dealing with diffusion and motion-induced diffraction [64]. The term also has meaning for memories designed to store subsequently arriving photons. Here the modes are temporal – the photons are distinguishable by when they are read in, but

otherwise identical. More simply put, in this latter sense this figure of merit answers the question of how many photons can be stored before some must be read out. In the context of long-distance quantum communication this capability eliminates the need to wait for the first photon to propagate to its target and for information to propagate back before the next photon is sent, thus increasing the rates. The alternative for single-mode memories is just to build more of them. To obscure the banality and brute-force nature of this solution, this plan is sometimes given the elaborate name *spatial multiplexing*, but this is not to say that the idea is bad. Indeed it is exceedingly rare for any kind of information technology to rely on higher dimensionality over duplicate elements, and if the memories are sufficiently cheap and simple, and useful enough to justify the investment, then history suggests engineering will eventually make this approach the most viable.


### Practical Measures of Noise and Output Photon Quality

For the storage of single photons or single-photon level coherent pulses the fidelity defined above is not directly evaluable. Instead, simpler and more specific measures that can be extracted from sets of storage and retrieval attempts can be used. Due to their significance to the experiments in this thesis I also include them here. The simplest of these proxy measures is the signal-to-noise ratio. Experiments are run with and without an input and the emitted numbers of photons are counted and compared:

$$\text{SNR} = \frac{N_{\text{out}} - N_{\text{noise}}}{N_{\text{noise}}}. \quad (1.10)$$

Generally the storage time is fixed for this kind of characterization, and the photons are counted within a window where the read-out is actually attempted. It is usually legitimate to widen or narrow this window as desired if consistency is maintained across all figures of merit calculated, for example if the end of the read-out contains much more noise than the beginning. In a practical application this kind of time-gating can be implemented on a hardware level. A similar parameter named the read-out noise floor, or  $\mu_1$ , is the number of input photons for which the output signal-to-noise ratio is one. If we were characterizing the memory with a perfect single photon input this is simply the inverse of the signal-to-noise ratio, but this is typically not the case. The noise floor parameter corrects for these source imperfections. For comparison between implementations  $\mu_1$  can be made independent of other figures by giving it as the ratio of the number of noise photons detected per retrieval attempt to the internal efficiency at zero storage time.

These measures concern the amount of noise, but not all noise is created equal. A property that quantifies the solitude of a single photon is a statistical measure called



**Figure 1.2:** To access the statistical nature of photons interference experiments must be performed. To determine the indistinguishability by Hong-Ou-Mandel interference the output of the memory is split on a 50:50 beam splitter  $BS_1$ , and coupled to fibers, where one path is considerably longer than the other to serve as a delay. Once a photon has been stored and read out, a second photon is stored for a time equal to the fiber delay. If the first photon took the long fiber to the detectors, and the second photon takes the short one, then both will arrive at the second 50:50 beam splitter  $BS_2$  at the same time. Coincidences in the clicks of the two photon counting detectors,  $D_1$  and  $D_2$ , will then yield insight into the photons behavior at the beam splitter. To determine the autocorrelation the first beam splitter and long fiber arm are omitted.

the autocorrelation. It represents the relative probability of multi-photon emission over time. An even more stringent measure is the indistinguishability between photons, an experimentalist's practical counterpart to the fidelity. Figure 1.2 illustrates how these properties can be measured with the Hanbury Brown and Twiss effect and Hong-Ou-Mandel interference respectively. For a single mode field the measured autocorrelation can now be defined as a ratio of photon detection probabilities [65],

$$g^{(2)} = \frac{p(D_1 \wedge D_2)}{p(D_1)p(D_2)}. \quad (1.11)$$

This formulation with detection probabilities can be converted to terms of field operators for theory, or estimated by measured photon detection rates. Let us consider some edge cases. If a photon is truly single, then it's impossible for both detectors to detect it and the numerator is zero resulting in  $g^{(2)} = 0$ . This effect is called antibunching and the resulting statistics are called sub-Poissonian. Coherent light sources emit photons independently of one another, which results in independent detection probabilities of the detectors yielding  $g^{(2)} = 1$ . This is the case of Poissonian statistics. Photons from thermal sources bunch, meaning that they are preferentially emitted together. This allows for higher values of  $g^{(2)}$ , up to 2 for classical light, bounded by a Cauchy-Schwartz inequality [66]. (A more precise analysis is deferred to chapter 4.)

Hong-Ou-Mandel interference [67] can be used to determine the indistinguishability of photons. This is a relative measure, so it cannot be determined by the storage and retrieval of individual photons alone. If two indistinguishable photons are in the modes



a and b when they interfere on the second beam splitter shown in figure 1.2, then they will both end up exiting together, either in mode c or in mode d. Symbolically this can be written with (mode occupation) number states as

$$|1, 1\rangle_{ab} \rightarrow \frac{|2, 0\rangle_{cd} - |0, 2\rangle_{cd}}{\sqrt{2}}. \quad (1.12)$$

The cross-correlation between the detectors measuring photons in the modes  $c$  and  $d$  can then be calculated to yield a coincidence probability  $p_{\text{coin}}$ , which relates to the indistinguishability or HOM visibility  $\mathcal{V}$  as

$$p_{\text{coin}} = \frac{1 - \mathcal{V}}{2}. \quad (1.13)$$

If the mode overlap can be treated as perfect, the visibility as solely impacted by the photon number, and the beam splitter inputs as symmetrical, then the relation between the visibility and autocorrelation is well approximated by  $\mathcal{V} = (1 + g^{(2)})^{-1}$  [68]. This estimate reveals the fairly intuitive classical limit of  $\mathcal{V} = 50\%$  corresponding to coherent light. To describe the real situations of arbitrary input fields and linear optical gates the indistinguishability can be modeled using the photons' spectral properties [69, 70]. Two kinds of indistinguishability are interesting here, on the one hand the indistinguishability of a photon emitted by a source directly and one that is stored, as well as the indistinguishability of two photons subsequently stored in the memory. The first is a relevant measure if the properties of the photon to be stored are particularly important, but the second is sufficient if memories are used to synchronize sources. More on this in section 2.3.3.



## Chapter 2

# Review of Memories for Light and Single Photon Sources

*If philosophy is a form of magic, then physics is a necromancy.*

Eugene Thacker

---


Having laid the ground work for evaluating the quality of memories we can now survey and compare implementations meaningfully. This chapter covers a variety of memories, starting with some basic reference points that are memories by our first definition – tools that solve the problem of information being lost to its user – but lack some typically desirable features. Then the focus falls on multi-level atomic systems storing light in matter. Such systems are intensively investigated in the context of the quantum science concerned with light-matter interaction generally, and for the storage of single photons in particular. I have organized the discussion by different memory schemes, but the figures of merit also vary prominently with the physical system used to implement them. To avoid too much repetition in the overview of protocols I have associated the discussion of different physical systems with specific schemes. Rare-earth ion doped crystals mainly feature in the section on echo protocols, whereas alkali atoms feature prominently thereafter, particularly in the section on electromagnetically induced transparency. This is a practical simplification, and I have included a few diverse examples in each section so as not to give too one-sided an impression of the research. In order to discuss the realization of interfaces in any depth, it is also necessary to understand the sources of the single photons. These are reviewed next, again beginning with figures of merit in section 2.2. Finally, I close the discussion with the envisioned applications in quantum communication and information processing and a brief look at their realizations so far.

## 2.1 Review of Memories for Light

Light's fast propagation and weak interactions, so useful in the transmission of information, are a double edged sword when information encoded in light is to be kept in place. Although not strictly necessary to constitute a memory, the properties of light are so extreme in this regard that the typical approach is to convert the light to an excitation in matter which is more inherently stationary. The simplest example is a two-level system connected by an electric dipole transition with which the light to be stored is resonant. Let us begin with the two-level system in the ground state, call it  $|g\rangle$ , and an incident photon. When the photon is absorbed, the two-level system ends up in the excited state,  $|e\rangle$ . For the lifetime of the excited state we can consider the photon to be stored. Unfortunately, by time-reversal symmetry, spontaneous emission retrieves this photon just as quickly as it was absorbed, and we lack any authority over the matter. What we require is control over the manner in which the photon is absorbed and emitted. Before we delve into how this control is established however, let us first examine the premise that a mapping to a state in matter is required to preserve the information contained in the light.

### 2.1.1 Delay Lines

If it is known in advance when a photon will arrive and when it will be needed, then it is generally much easier to delay it by that amount of time than it is to write it into a medium and read it out on demand. The photon is allowed to propagate locally for a fixed distance  $L$ , either in free space or in a medium, delaying it by  $\tau = \frac{L}{c}$  or  $\tau = \frac{nL}{c}$  respectively, where  $c$  is the speed of light and  $n$  is the refractive index of the medium. Propagation in free space or fibers provides this delay, but because the speed of light is high, long delays require long distances. Here exponentially scaling losses begin playing an important role. Let us examine the fiber based scenario more closely to get our bearings. Assuming a refractive index of  $n_f = 1.5$  for a glass fiber, we are able to get a  $5 \mu\text{s km}^{-1}$  delay by coupling our signal into it, a process that will also involve some loss before any delay has been achieved. Inside the fiber, propagation losses are dominated by Rayleigh scattering in the visible [71], and absorption from lines in silica itself starting towards the end of the near infrared [72]. The magnitude of losses varies widely with wavelength, with a minimum generally occurring for light in the telecom C band. An overview of common values and records can be found in [73]. Losses are usually given as the (decibel) attenuation of the intensity per distance  $L$ , with  $aL = 10 \log_{10} \left( \frac{I_{\text{in}}}{I_{\text{out}}} \right)$ . This translates into percentile efficiency as  $\eta = 10^{-aL/10}$ . Dispersion, especially relevant for ultrashort pulses, places upper bounds on the fidelity of this delay. Indeed, a delay of this kind can be analyzed exactly like a memory, the



**Figure 2.1:** A scheme to delay a photon or light pulse by an integer factor of the single loop round-trip time. A Pockels cell acting as a fast polarization switch effectively turns a polarizing beam splitter into a variable delay line.

only caveat being the fixed nature of the storage time. The delay can be increased more artfully, and possibly more efficiently, by using a medium that slows down the group velocity of the light [74]. This can be done, for instance, with electromagnetically induced transparency in atomic vapors, which can provide large fractional delays (delays of many optical pulse widths) for very short pulses with little distortion [75]. This gets straight into the same physical realm we will encounter again when the light is to be stored in the medium, rather than just slowed by it, so I will elaborate on it further in section 2.1.4.

It is technically simple to extend the hold time achieved by a delay line by an integer factor through looping it. The storage and manipulation of streams of optical pulses in loops has been investigated and found to be feasible since the early '90s [76, 77]. For a simple example of how light can be trapped in and retrieved from a loop, consider the scheme drawn in figure 2.1. Light is initially horizontally polarized and transmitted through a polarizing beam splitter. The light is allowed to propagate for some distance, or is additionally delayed by any desired method, before looping back around after the single loop round-trip time  $\tau$ . Before it impinges on the beam splitter a second time, a voltage pulse is applied to a Pockels cell to rotate the light's polarization into the vertical. This causes it to be reflected at the beam splitter, which keeps it in the loop. After  $N$  round-trips the Pockels cell is pulsed again to rotate the polarization once more, now back to the horizontal, and the light exits the loop at the second beam splitter port having been delayed by  $N\tau$ . A scheme like this is a sufficient memory, for instance, to synchronize pulsed spontaneous photon sources as implemented by Kaneda et al. [78], and is thus (correctly) called a memory therein without any mapping to matter having occurred. An optical loop has also been interfaced with an atomic photon source to enable the heralding and manipulation of a chain of subsequently emitted photons [79]. We will return to synchronization applications in section 2.3.3. Simple expansions and alternative electro-optic modulations are also possible, for example to delay not just a photon initially horizontally polarized but one of arbitrary polarization acting as a polarization qubit [80]. Cavities with switchable mirror separations are another

method to control the input and output of a delay line [81].

### 2.1.2 Optimal Cloning

Back in section 1.2.1 I noted that an unknown quantum state cannot be perfectly duplicated. Indeed, this is the underlying reason quantum cryptography can provide physical security against eavesdroppers in the first place [82]. It does, however, also beg the question of whether at least a reasonable copy can be made. Answering in the affirmative, optimal cloning devices were described first for single discrete states [83] then generalized to copying  $N$  systems to  $M$  systems [84]. Upper bounds on the fidelity of these processes were found based on foundational concepts like forbidding faster-than-light signaling to enforce consistency with relativity [85] or limits on measurement<sup>1</sup> [86], and the fidelity of the optimal processes was found to be bounded by  $\mathcal{F} \leq \frac{2M+1}{3M}$ . Experimentally, cloning machines for single photons realizing the mapping  $1 \rightarrow 2$  through spontaneous parametric downconversion [87] and back-to-back results achieving the cloning by exploiting photon bunching on a beam splitter [88, 89] have all reached fidelities very near the optimal  $\mathcal{F} = \frac{5}{6}$ . Continuous variable states, e.g. coherent states of light, have also been investigated [90]. In this case the upper bound on the fidelity is  $\mathcal{F} \leq \frac{M}{2M-1}$  for  $M$  copies of one state [91]. Since the first descriptions and implementations, concrete proposals have been made for larger scale optimal cloning. These include, for example, one scheme to amplify photons to levels detectable by eye while preserving some of their correlations [92].

Compared to memory schemes, quantum cloning is easily implemented; beam splitters and parametric amplifiers are enough to clone optimally [93]. Together with variable delays as described above, they set thresholds in the figures of merit for state-mapping memories to beat to prove that they are worthwhile. This is more than an idle classification exercise too, as cloning machines are tools we must assume an eavesdropper or untrustworthy party in a quantum communication protocol has access too. This means these bounds directly relate to the practical security of these protocols.

### 2.1.3 Approaches to Optical Depth

Baselines established, we can now turn to schemes designed to convert a photon into an excitation in matter in a reversible way. First and foremost there is the general question of how to achieve sufficiently strong coupling between light and matter, effectively high optical depths, so that storage can principally be efficient. Practically, there is some variation in approaches, but a single three-level atom in free space generally does not couple strongly enough to an incident photon for reliable storage. Some improvement

---

<sup>1</sup>In this context quantum cloning machines can be interpreted as devices that translate quantum information into classical information.

has been shown to be possible with tight focusing, and nearly 10% absorption of a probe by a single atom in free space has been reported [94]. The idea of this confocal microscope approach is to overlap the incoming optical mode with that of the atom's spontaneous emission. The results are better than predicted by simple models [95], but accurately reproduced when the high power lens is treated without too many simplifying assumptions, particularly with regards to the influence of the lens on the light's polarization [96]. Unfortunately, obtaining diffraction limited performance for such tight foci and placing the lenses sufficiently precisely has proven to be no less challenging technically than long established cavity based approaches, and the full theoretical pay-offs of near unity absorption without cavity associated losses [97] have so far remained theoretical. Thus, for now at least, the main approaches for increasing the coupling remain placing atoms in high finesse cavities, using ensembles of atoms, or combinations of the two.


As cavities are a great aid in establishing quantum control over light and matter, cavity-based quantum electrodynamics (cQED) was fertile ground for exciting proposals and experiments early on. Note also that the theoretical description of a single optical mode in a cavity is considerably simpler than treating the continuum of radiation modes in free space [98]. By the turn of the millennium, experimental demonstrations included entanglement generation for up to three parties, phase gating, and absorption-free photon detection [99]. One of the earliest proposals for quantum state transfer and entanglement distribution in a quantum network relied on states coupled by a Raman transition in cavities [100]. Implementations for weak coherent states [101] and polarization qubits based on them [102] were forthcoming in the following decade, then scaled to elementary links with single photons in 2012 [103]. Indeed, I think the experiment by Ritter et al. [103] is a good example as it typifies some features of memories using single cold atoms in cavities. There are two nodes, A and B, consisting of single atoms in cavities. At node A some quantum state is encoded in the atom using its Zeeman levels. Then this state is mapped to the polarization of a photon, which is sent over to B where it is mapped back to an atom recreating the state originally encoded at A. The fidelity is very close to unity and the principally achievable efficiency is reasonable, although in practice some of it is deliberately traded-off to minimize noise. On the other hand, the bandwidth is just a few MHz and impossible to increase without sacrificing the performance due to the required cavity finesse, which is given as  $6 \times 10^4$ . Due to the preparation involved in optically pumping and cooling the atom, the maximum possible repetition rate without including a storage time is 10 kHz. The protocol works really well because the control over the interaction is extraordinarily good, but the apparatus is necessarily slow and complex. The latter remark is not criticism, but it places the use case of networks made with such nodes in fundamental research, where the control, efficiency, and fidelity must be maximal

to properly investigate emergent phenomenon like quantum phase transitions of light [104] and the percolation of entanglement [105]. This is explicitly also the conclusion of Ritter et al. as these are their examples.

What about memories aiming to function not as tools for fundamental research, but in the more applied sense, like to enable quantum communication or other applications discussed in section 2.3? Lower finesse cavities which limit the bandwidth far less can be extremely useful when the optical depth of the medium cannot easily be enhanced by other means, such as in the case of rare-earth ion doped crystals. When the cavity and medium are impedance-matched [106, 107], meaning that the absorption by the medium per round-trip equals the transmission of the coupling mirror, then the signal can be absorbed completely even if the storage medium only has an optical depth of order unity. Implementations of this idea in memories based on  $\text{Pr}^{3+}$  and  $\text{Eu}^{3+}$  doped  $\text{Y}_2\text{SiO}_5$ , weakly absorbing crystals on their own, each achieved total internal efficiencies  $\eta > 50\%$  [108, 109]. An example given in [106] quantifies the bandwidth trade-off for these kinds of impedance-matched cavities, calculating a cavity bandwidth of 480 MHz at a finesse of 31 for a 1 cm long crystal with a peak optical depth of 1 reaching 92% efficiency. Clearly this is a different regime than that applicable to high-finesse cavities. A cavity of moderate finesse placed around the storage medium, explicitly in a regime remaining compatible with a high memory acceptance bandwidth, has also been proposed as a four-wave mixing noise suppression mechanism in Raman memory [110]. The enhancement is secondary here as the memory discussed is in hot vapor and thus optical depth is readily available. The cavity is tuned such that the signal is resonant while the noise is anti-resonant. Again the practical situation is quite different than for cQED with single atoms as far as bandwidth limitations imposed by the cavity stand. As a final note on cavities, the inverse relation between acceptance bandwidth and cavity enhancement can be circumvented by the technique of white-light cavities [111, 112], originally having gravitational wave detection in mind. The idea is to cancel the phase shift experienced by off-resonant frequencies in the cavity using a medium with negative dispersion in the cavity, for instance a three-level system. Implementing this and memory operation is the topic of a very recent proposal promising to achieve high efficiencies, long storage times, and large bandwidth simultaneously [113].

Another approach to enhancing the coupling is through ensembles. Although the interaction probability with any given atom remains small, it asymptotically approaches one with increasing atom number as the optical depth becomes much larger than unity. This approach is the focus of this thesis, so rather than discussing it abstractly I will implicitly assume it as we delve into the details about specific schemes. With regards to the interaction strength most systems in the solid state require help with their optical depths, but hot and cold atomic ensembles are typically fine without cavities. Besides the considerations discussed in this section however, memory protocols are generally





**Figure 2.2:** A basic lambda scheme consisting of the metastable ground states  $|g\rangle$  and  $|s\rangle$ , each connected to the excited state  $|e\rangle$  by dipole transitions. The levels are labeled such that the signal  $\mathcal{E}$  is (near) resonant to  $|g\rangle \rightarrow |e\rangle$  and the control  $\Omega$  is (near) resonant to  $|s\rangle \rightarrow |e\rangle$ , with the possibility of a two-photon detuning  $\Delta$  from the excited state equal for both transitions.

agnostic to the method by which the interaction is enhanced [114].

#### 2.1.4 Lambda Memories

As we immediately realized when contemplating storing a photon in a two-level system, a little more complexity of the matter system is desirable for useful quantum control. A natural solution is the addition of a third level  $|s\rangle$ . A popular choice of this third level is a metastable ground state, and due to the shape of this level scheme, sketched in figure 2.2, this case is named after the letter  $\Lambda$ . First, the system is prepared in the state  $|g\rangle$ . From the perspective of the photon to be stored, the state  $|s\rangle$  is dark, i.e. the coupling between them is negligible. When the photon is absorbed on the transition  $|g\rangle \rightarrow |e\rangle$ , the control  $\Omega$  maps the excitation to the metastable ground state  $|s\rangle$ , which is comparatively long lived and rarely limited by spontaneous emission. For ensembles of  $N$  atoms the resulting state is a spin wave of the ground states, explicitly a collective excitation of the form

$$\frac{1}{\sqrt{N}} \sum_i |g\rangle_i \langle s|. \quad (2.1)$$

Continuing our simplified spontaneous emission estimation, Fermi's golden rule yields the decay rates from  $|e\rangle$  to  $|g\rangle$  and from  $|s\rangle$  to  $|g\rangle$  as  $\Gamma \propto \omega^3 |\mu|^2$ , where  $\omega$  is the frequency of the respective transition and  $\mu$  is its dipole moment. For common choices of these types of systems the difference is immense, for example consider the hyperfine split ground states and D line excited states in alkalis where there are around 5 orders of magnitude difference in the transition frequencies [115], which are then cubed in

the decay rates. Now, often the lifetimes of these hyperfine ground states are limited by something else, for example collisions, but the difference to the excited states is generally still best expressed in orders of magnitude. After the excitation is transferred to  $|s\rangle$ , the memory can be read out by applying a control pulse in the reverse direction. The read-out process then is the time-reversal of the read-in and guaranteed to be maximally effective. This is called *backwards retrieval*. Often, either for geometrical reasons of the physical setup like difficulties in collecting the read out signal effectively this way, or due to mode or phase matching considerations, it is preferable to read the memory out in the forward direction. In this case the photon continues on in the direction it was going during read-in. As this is not a perfect time reversal of the read-in the efficiencies tend to scale slightly worse for forwards retrieval, and in some cases do not tend towards unity with the optical depth.

There are a plethora of memory protocols sharing this basic commonality of the  $\Lambda$ -scheme. These differ widely in the process responsible for the mapping, as the detuning from the excited state moves the system through different scattering regimes, or the nature of the control, which need not even be optical, changes the underlying physics. Fortuitously, the storage and retrieval of photons in all of these systems can be described by a universal approach agnostic to the exact methods [114]. This generic analysis yields two famous results applicable to all implementations. Firstly, the maximum achievable efficiency of the storage and retrieval process depends exclusively on the optical depth of the storage medium. Secondly, optimal storage and retrieval are time-reversals of each other. In cavities a minor modification is required for the first of these claims: the efficiency is determined by the cavity cooperativity parameter  $C$ , with  $\eta = \frac{2C}{2C+1}$  optimally.<sup>2</sup> Here  $C$  is just the optical depth of the atom or ensemble in the cavity scaled by the number of passes a photon makes before leaking out [116]. In terms of atom number  $N$ , the single pass efficiency scales as  $\eta \propto 1 - \frac{1}{\sqrt{N}}$ , while in cavities it scales as  $\eta \propto 1 - \frac{1}{N}$  [121]. The main conclusions that can be drawn from this model is that the maximal efficiencies across implementations are the same, and

---


<sup>2</sup>In their series of papers on photon storage in generic  $\Lambda$ -systems [114, 116, 117, 118, 119] Gorshkov et al. use the optical depth definition conventional for theorists:  $d = \frac{g^2 NL}{\gamma_c}$ , where  $g$  is the single-photon coupling constant of the transition and  $\gamma$  is the coherence decay rate of the excited state (with  $2\gamma = \Gamma$  for purely radiative decay). This results in it being half as large as the conventional optical depth defined in experimental AMO physics as  $OD = \sigma nL = \log(I_{\text{in}}/I_{\text{out}})$ , where  $\sigma$  is the atomic absorption cross-section on resonance and  $n$  is the number density. When quoting their results, I have modified them to use OD to aid comparability to experimentally measured optical depths. The same is true for the cooperativity  $C$ , but as we will not be seeing it again I have refrained from renaming it. Unfortunately the translation between theoretical figures and measurements in this context is not finished with that, but turns out to be a recurring stumbling block due to a multitude of linewidth and transition strength scaling preferences. In his report [120] on simulations implementing this model, testing the practical feasibility of storing quantum dot photons in cold rubidium ensembles and laying some of the original motivating groundwork for quantum memory projects in Basel, Matt Rakher accounts for and illuminates these differences in an exemplary and helpful way.

that given sufficient optical depth they are achieved by optimizing the control to the signal to be stored (or vice-versa). This is not to minimize the distinctions between physically different  $\Lambda$ -memories. Other figures of merit and technical feasibility, both generally and of achieving this high efficiency specifically, all vary widely. Indeed, let us now examine some of these protocols in detail.

### Photon Echo

First observed with spins in the context of nuclear magnetic resonance [122], and later with light [123], echoes occur when a signal is absorbed on an inhomogeneously broadened transition, upon which the induced coherence dephases. After a time  $\tau$ , the sign of the dephasing process is reversed, and by time-reversal symmetry the coherence is restored to the original state after another delay of  $\tau$ , causing the signal to be emitted once more. In principle, no third atomic level is necessary. An ensemble of two-level atoms, or more typically ions, with the states  $|g\rangle$  and  $|e\rangle$  is prepared in a coherent superposition by a  $\frac{\pi}{2}$  pulse, the write pulse. Then the signal, which may consist of multiple pulses or photons, is sent to the atoms. As long as all this happens faster than the homogeneous relaxation time of the excited state, the signal interacts coherently creating a population grating, that is it imprints its frequency spectrum on the atoms as a frequency-dependent modulation on the absorption profile. In other words, the spectral profile of the excited atoms now reflects the power spectrum of the signal [124]. Again, the coherence dephases for  $\tau$  until a  $\pi$  pulse, the read pulse, is applied to invert the populations leading to their rephasing and emission of the signal after another delay  $\tau$ . Manipulating the write and read pulses makes it possible to perform arbitrary operations on the data [125]. A simple example of such an operation is temporally reversing the signal upon read-out by reversing the order of the write pulse and the signal. Investigations into using such systems as memories began with classical light [126] and were originally motivated by the potentially high bit densities, and of course the infinitely alluring possibility of producing “ultra-high-speed holographic motion pictures” [127].

Serious application of echo techniques to the quantum regime with quantum communication applications in mind needed to solve some problems of the classical implementations first. On the one hand, the classical protocol achieved only small efficiencies, on the order of a few percent. Either the optical depth was low and the read-in efficiency was diminished, or it was high and the signal tended to be reabsorbed upon read-out. On the other hand, the classical two-pulse echo technique suffered not only from fluorescence noise, but it was also eventually shown that the rephasing pulse necessarily stretches as it propagates through the sample, producing a trailing tail that can not be separated from the signal echo [128]. The two-level atom approach was



**Figure 2.3:** Storage (left) and retrieval (right) of a weak signal or photon  $\mathcal{E}$  in the CRIB memory scheme. The excited state is inhomogeneously broadened and this broadening must be reversed to perform the read-out, which is represented by the reversed color gradients of the excited line shape. The protocol is performed on resonance and the control pulse  $\Omega$  is a  $\pi$ -pulse between the states  $|e\rangle$  and  $|s\rangle$ .

abandoned in the proposals of the two most prominent modern schemes, controlled reversible inhomogeneous broadening (CRIB) [129] and atomic frequency combs (AFC) [130], in favor of lambda systems. Let us now look at how these schemes work more closely.

The level scheme envisioned in CRIB is sketched in figure 2.3, maintaining the previously established labeling of the states. After preparation of the atomic ensemble in the ground state  $|g\rangle$ , a weak signal or photon  $\mathcal{E}$  is absorbed on the inhomogeneously broadened transition  $|g\rangle \rightarrow |e\rangle$ , which must be sufficiently wide spectrally to absorb all the signal's spectral components, as well as optically thick. Immediately afterwards a  $\pi$ -pulse  $\Omega$  maps the excitation to the metastable ground state  $|s\rangle$  for long storage. Reversing the broadening now enables unity read-out efficiency by time-reversal symmetry. In the original proposal by Moiseev and Kröll [129] they take the storage medium to be a Doppler-broadened atom gas. To reverse the sign of the broadening for their read-out they use an old trick of atomic physics famous due to its utility in Doppler-free spectroscopy, they apply the read-out  $\pi$ -pulse in the counter-propagating direction of the read-in pulse.

The original CRIB scheme is limited by atomic motion, either out of the interaction region, i.e. the spatial mode addressed by the control laser, or due to collisions changing the atoms' velocities and phases. This motivated modifications which moved the scheme into rare-earth ion doped crystals both conceptually [131, 132] and practically [133], for instance  $\text{Eu}^{3+}$ ,  $\text{Pr}^{3+}$ , or  $\text{Er}^{3+}$  doped  $\text{Y}_2\text{SiO}_5$ ,  $\text{CaF}_2$ , or  $\text{LaF}_3$  to name a few.<sup>3</sup> The review by Thiel et al. [136] contains a long table collecting the key properties of various rare-earth doped materials. At cryogenic temperatures the ions move so little


<sup>3</sup>As I disclaimed at the beginning of this chapter, the following focus on rare-earth ion doped crystals as platforms for photon echo memories mostly serves to divide and simplify the discussion. CRIB memories implemented in hot vapor [134] or cold atomic ensembles [135] are completely viable and have achieved impressive fidelities and efficiencies.

that their coherence is limited by spin-spin interactions rather than by motional or collisional losses. Hyperfine coherence times of milliseconds are typical without further intervention, but around the time of these proposals dynamic decoupling techniques originally developed for NMR spectroscopy were already being used to extend these lifetimes to seconds [137]. Later a record of 6 hours was set in a system with optically addressable nuclear spins using an applied B-field to eliminate first-order sensitivity to field fluctuations<sup>4</sup> [139]. This longevity is due to the investigated transitions occurring within the partially filled 4f electron shells, which are shielded from the environment by filled 5s<sup>2</sup> and 5p<sup>6</sup> shells, which means that even the excited states can have lifetimes of milliseconds [136]. Indeed, there are proposals [140] and experiments [141] which forgo the lambda approach and forge ahead with 2-level atoms. The motivation for foregoing storage in a ground state is usually that the required optical control pulses produce noise, and while the storage times are necessarily shorter, it is possible to read out storage in an excited state on demand without optical control [142]. Either way, the exchange of the typical physical system to ions for echo memories had already taken place when the atomic frequency comb memory scheme was proposed, so these remarks apply to it as well.

Due to variations in the host material, that is the fixed but unique environment of each ion across the length of the crystal, the absorption lines are initially broadened inhomogeneously from their natural homogeneous linewidths on the order of kilohertz to gigahertz or more. The difference in linewidths can reach 8 orders of magnitude, which has occasionally been styled as the number of addressable frequency channels. However this broadening is not of the reversible type, so to address an isolated system without inducing noise from other lines it is removed by a combination of optical pumping and hole-burning [143]. One possible approach in particular is to first depletion pump a wide spectral region of ions using chirped pulses, to then refill a narrow line by pumping back those hyperfine states experiencing a specific shift. Once the system is prepared in this way a controllable inhomogeneous broadening of the line must be introduced. This is achieved by applying a varying electric field to cause a Stark shift. In the proposal by Nilsson and Kröll [131] the field is varied along the propagation direction, a scheme later named longitudinal-CRIB (lCRIB) or gradient echo memory (GEM), but the transverse direction can also be used [144] yielding transverse-CRIB (tCRIB). In the former case the absorption line of each ion at position  $z$  is narrow, but varies through the medium with  $\Delta = \chi z$ , whereas in the latter case each absorption line is identically broadened independently of the position. The induced broadening can now be reversed

---

<sup>4</sup>In atomic physics this technique is famous due to atomic clocks, and transitions made insensitive to fluctuations of the magnetic field in this way are thus called clock transitions. This nomenclature was not adopted when the technique was pioneered in the solid state [138], and while it is sometimes used, the slightly less pithy name “ZEFOZ transitions” (zero first order Zeeman) is more common in these systems.



**Figure 2.4:** After initial preparation in the state  $|g\rangle$ , a frequency selective pumping laser  $P_\nu$  selectively transfers atoms to an auxiliary state  $|\text{aux}\rangle$ . The desired comb-like absorption profile consists of peaks with spectral width  $\gamma$ , evenly spaced with a separation  $\Delta$ . The finesse of the comb structure is defined as  $F = \Delta/\gamma$ . The comb absorbs the signal  $\mathcal{E}$  and begins dephasing, but due to the periodic structure of the absorption spectrum the comb excitation rephases on its own after a time  $2\pi/\Delta$ . Intermittently the excitation can be transferred to a long lived spin wave in  $|s\rangle$  as in CRIB.

by inverting the polarity of the applied electric field. The efficiency of lCRIB is the same for backwards and forwards retrieval and scales as  $\eta(t) \propto (1 - \exp(-\text{OD}_{\text{eff}}))^2 f(t)$ , where  $\text{OD}_{\text{eff}} = \text{OD} \frac{\gamma_0}{\gamma}$  is the effective optical depth. In words, the optical depth of the initial line is scaled by the ratio of the initial linewidth  $\gamma_0$  to the inhomogeneously broadened linewidth  $\gamma$  [55]. The function  $f(t)$  is intended to generically represent the decoherence with time of the storage state, for example if storing in the excited state we would have  $f(t) = \text{sinc}^2(\gamma_0 t)$ . For tCRIB the retrieval direction matters, backwards retrieval scaling as  $\eta \propto (1 - \exp(-\text{OD}))$  and forwards retrieval as  $\eta \propto \text{OD}^2 \exp(-\text{OD})$  with the optical depth often given as the product of the crystal's absorption coefficient and its length  $\text{OD} = \alpha L$  [125]. Note that in the final case the maximum value is not 1 but  $\frac{4}{e^2} \approx 54\%$  [144]. Analytical solutions for the fidelity have also been derived [145]. The conclusions drawn from these analyses are that high fidelities can be reached, but that there are optimal parameters of the interaction time, optical depth, and inhomogeneous broadening that maximize it. Other figures of merit of this scheme also depend on the materials used, in particular the possible bandwidth will be determined by how close the next nearest hyperfine transitions lie to the isolated  $\Lambda$ -scheme. For instance, analyzing the case of  $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ , Nilsson and Kröll [131] estimate a maximal bandwidth of 4 MHz.

Multi-mode capacity can be analyzed in a universal way using the Schmidt decomposition of the Green's function of the memory interaction [146]. For both field directions in CRIB the capacity scales linearly in the optical depth. This is better

than the scaling with the root of the optical depth in adiabatic memory protocols, like electromagnetically induced transparency and Raman protocols, without an applied inhomogeneous broadening. The greatest advantage in this regard is however found in atomic frequency comb memories, where the number of potential modes is independent of the optical depth and proportional to the number of peaks in the comb. This was a significant part of the motivation in the original AFC proposal [130] and proof of principle [147]. Notably it was thereby the only scheme that could plausibly physically implement the conditions [148] for reasonably fast entanglement distribution via a quantum repeater at the time. The AFC scheme modifies CRIB by tailoring the inhomogeneously broadened absorbing line into a comb structure by pumping to an auxiliary level as illustrated in figure 2.4. The signal to be stored must be spectrally broader than  $\Delta$  and narrower than the comb's total width  $\Gamma$ , which is itself limited physically by the frequency splitting of the ground states. Under this condition the energy-time Heisenberg uncertainty ensures spectral averaging of the comb structure that leads to uniform absorption over the signal bandwidth [130].

After absorption of a single spatial mode with wave number  $k$ , by  $N$  atoms at the positions  $z_j$  and with detunings  $\delta_j$  with respect to the laser frequency, each contributing an amplitude  $c_j$  to the absorption, the system is in a Dicke state between the ground and excited states of the form

$$|\psi\rangle = \sum_{j=1}^N c_j e^{i\delta_j t} e^{-ikz_j} |g_1 \dots e_j \dots g_N\rangle. \quad (2.2)$$

By design, the phase term  $e^{i\delta_j t}$  causes rapid dephasing, but the trick of the discrete absorption spectrum ensures spontaneous rephasing after a time  $2\pi/\Delta$ . The mapping to  $|s\rangle$  by a  $\pi$ -pulse control field for longer storage locks in these phases as there is no comb structure on this transition. The efficiency of this protocol, initially calculated assuming Gaussian comb teeth, is  $\eta \approx (1 - \exp(-\text{OD}_{\text{eff}}))^2 \exp(-7/F^2)$ , where the effective optical depth is now  $\text{OD}_{\text{eff}} \approx \text{OD}/F$  [55]. Equations for Lorentzian as well as arbitrary shapes of the comb teeth can be found in [149]. Unsurprisingly and independently of the exact comb shape, we learn that the removal of atoms from contributing to the absorption by creating the comb diminishes the optical depth available, but the final term also implies that the comb finesse must be high for efficient storage, reinforcing the requirements on high optical depth even more.

Although I initially diminished the importance of the operating wavelength of memories due to the growing effectiveness of frequency conversion, it is worth mentioning that erbium doped crystals have transitions in the telecom C-band. Memories using this medium would be directly compatible with long-distance fiber channels without frequency conversion, simplifying network integration. Without going into further

detail, the situation is complicated somewhat because erbium ions are Kramers ions, that is they have an odd number of electrons yielding half-integer spin. This means that the ZEFOZ technique used to extend the coherence times discussed above do not work – Kramers ions do not have optical clock transitions. That said, large magnetic fields (7 T) have been found to produce coherence times around 1 s for these types of ions as well [150]. Whether this approach significantly improves upon the experimental complexity versus frequency conversion is, in my opinion, debatable. Indeed, now having covered the figures of merit achieved by rare-earth ion doped crystal based echo memories it must be said that their experimental overhead is among the greatest across implementations. A cryostat, applied magnetic fields – which must either be large (Kramers ions case) or precise (ZEFOZ case) – applied electric fields, comparatively elaborate optical state preparation and most likely impedance-matched cavities are at serious odds with miniaturization, field-operability, and turn-key functionality for a lay end-user. Even in an optimistic vision of a realized quantum internet these memories are infrastructure set-pieces at the heart of sporadically situated repeater stations, maintained and operated by expert technicians. This is not a final verdict on the role of echo protocols *per se*. Recently an AFC memory was demonstrated in hot atoms wherein the comb was prepared using velocity-selective pumping [151], and other methods for preparing combs in hot atoms have been proposed [152]. Time will tell whether these implementations in technically simpler systems go beyond the proof-of-principle stage and achieve competitive figures of merit.

### Electromagnetically Induced Transparency

Electromagnetically induced transparency, or EIT, is an effect where a transition in an optically thick medium becomes transparent to light resonant with it [153]. Somewhat surprisingly considering the amount of quantum optics that had already been done at the point, its first description [154] and observation [155] date to the early '90s. To understand how light is stored by EIT, let us first consider what happens when a control laser  $\Omega$ , named thus after its Rabi frequency, is continuously shone on the transition  $|s\rangle \rightarrow |e\rangle$  in our standard lambda-system test case. The regimes of the resulting effect, called static EIT, are delineated by the intensity of the control, and are shown in figure 2.5. The laser modifies the atomic response due to quantum interference between the absorption pathways  $|g\rangle \rightarrow |e\rangle$  and  $|g\rangle \rightarrow |e\rangle \rightarrow |s\rangle \rightarrow |e\rangle$  [156]. A complete picture of the exact behavior of the susceptibility of the medium under these conditions can be found in reviews of EIT, e.g. equation 13 in [156], or equation 21 therein for a form that emphasizes the role of quantum interference. The resulting refractive index  $n = \Re\sqrt{1 + \chi}$  and extinction coefficient  $\kappa = \Im\sqrt{1 + \chi}$  are plotted schematically on the upper right side of figure 2.5. The significance of these plots lies in the relation



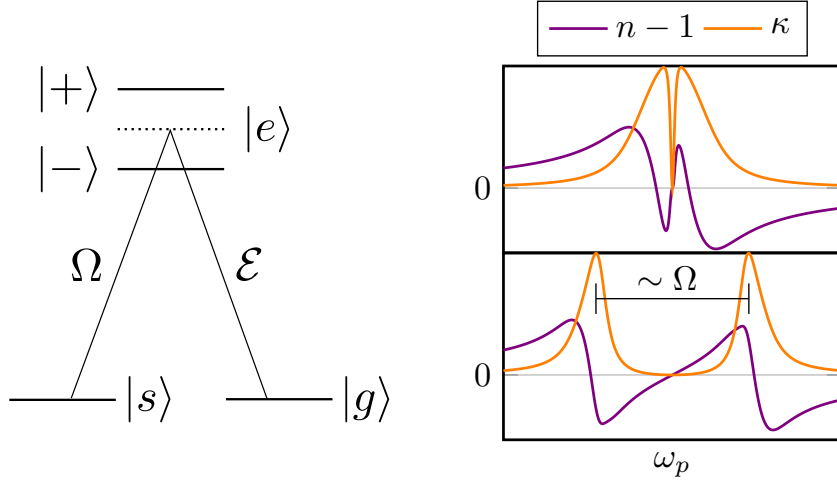
between the dispersion relation of a medium and the group velocity of light through it, with  $v_g \propto \left(\frac{\partial n}{\partial \omega}\right)^{-1}$  [157]. In other words, static EIT can reduce the group velocity of light, and indeed extremely leisurely light velocities of mere meters per second have thus been produced [158]. Although not directly relevant further, note that the full implications of  $v_g \propto \left(\frac{\partial n}{\partial \omega}\right)^{-1}$  include the possibility of light with  $v_g > c$ , as well as  $v_g < 0$ , which describes a scenario wherein the peak of a light pulse exits a medium before it enters it. For a broader introduction to slow and fast light, as well as some historical context – indeed observations of light pulses moving at unusual speeds far predate the description of EIT – I recommend the chapter by Boyd and Gauthier [159].

This interference based description is necessary when  $\Omega$  is smaller than the excited state linewidth. Transparency is induced on resonance, but there is not yet a significant splitting in the absorption line. This is the typical EIT parameter regime. For larger  $\Omega$ , however, a simpler picture presents itself. Consider that the electric dipole interaction Hamiltonian  $\mathcal{H} = \frac{\hbar\Omega}{2} |e\rangle\langle s| + \text{h.c.}$  is diagonal in the dressed states  $\mathcal{H} = \frac{\hbar\Omega}{2} (|+\rangle\langle +| - |-\rangle\langle -|)$ , with  $\sqrt{2}|\pm\rangle = |e\rangle \pm |s\rangle$ . The resulting double resonance at energies shifted by  $\pm \frac{\hbar\Omega}{2}$  from the natural level, known as an Autler-Townes doublet, is illustrated on the left side of figure 2.5. The refractive index and extinction coefficient in this case are plotted on the lower right. In this, the Autler-Townes regime, the energy splitting of the levels  $\hbar\Omega$  dominates the size of the transparency window. Varying  $\Omega$  tunes continuously between the regimes of transparency and line-splitting. This kind of dressed doublet shows up in all kinds of research fields, from optical spectroscopy, over cavity QED, to laser cooling [160].

So far it appears that what EIT produces is a kind of delay line, and with that it already has plenty of possible applications in optical signal processing including modulation and switching [161]. How exactly the effect can be used to store light was worked out around a decade after its first observation [162], but our simple picture already provides some clues. The width of the induced transparency window is determined by the control Rabi frequency – for large  $\Omega$  they are approximately equal. A signal now enters the medium and, assuming that it is not too long (quantified momentarily), compresses to fit within its physical extent as it slows down. If the control is now attenuated in an approach called dynamic EIT, then the window narrows and the slope in  $n$  grows steadily steeper until it diverges as  $\Omega$  falls to zero. This stops the light. Turning the control back on at a later time reverses this process and liberates the signal to move once more. In truth a description that properly explains why the signal is at no point absorbed<sup>5</sup> requires a full quantum treatment with light-matter quasi-particles named dark-state polaritons, which are superpositions of photonic and spin wave components [163]. This treatment also yields the condition that the control is

---

<sup>5</sup>This is decidedly not obvious in the description I have presented, after all as the transparency window is shut the absorbing peaks in  $\kappa$  also close in on the signal.



**Figure 2.5:** *left* In static EIT a strong control laser with Rabi frequency  $\Omega$  continuously dresses the transition  $|s\rangle \rightarrow |e\rangle$  inducing transparency and producing an Autler-Townes doublet for sufficiently intense illumination. The behavior depends on how the Rabi frequency of the control compares to the linewidth of  $|e\rangle$ . *right* A weak probe  $\mathcal{E}$  with the frequency  $\omega_p$  experiences the medium's rapidly changing dispersion due to the presence of this control. The refractive index and extinction coefficient are plotted schematically following [156] for the EIT (upper plot) and Autler-Townes (lower plot) regimes of  $\Omega$  to illustrate what happens to this probe.

switched smoothly (adiabatically) [164, 165] to ensure loss free and coherent transfer of a photon to an atomic state.

This final condition of adiabaticity is key to the analytical solutions for optimal control [117], but not actually necessary for highly efficient storage. Numerical optimization of the control with gradient ascent methods [119] also yields high efficiency when this condition is not met, as might be the case when input signals are very short. The general lambda-scheme result of efficiency depending on the optical depth benefits from a particularly intuitive explanation in the case of EIT memories. Let us consider the initial conditions that must hold for the dynamic EIT scheme to work. For a storage medium of length  $L$  and a signal of duration  $t$ , it must be the case that the spatial extent of the probe  $L_p = tv_g \leq L$  if the entirety of the signal is to fit inside the storage medium at once [166]. A related concern is whether the signal's spectrum fits inside the transparency window. The spectral full width at half maximum of the transparency window opened in the EIT regime with an optically dense medium, derived from the susceptibility in [167], is

$$\delta\omega_{\text{EIT}} = \frac{\sqrt{2}|\Omega|^2}{4\Gamma\sqrt{\text{OD}}}. \quad (2.3)$$

Here  $\Gamma$  is the radiative decay rate of the excited state.<sup>6</sup> We can thus require  $t^{-1} \leq \delta\omega_{\text{EIT}}$ . Still in the static EIT case, the relevant comparative measure to the signal duration is the group delay of the medium, which relates as  $\tau = \sqrt{\text{OD}}/\delta\omega_{\text{EIT}}$  [166]. We can now summarize these conditions in relative terms in the inequalities  $\frac{L}{L_p} \sim \frac{\tau}{t} \leq \sqrt{\text{OD}}$ . The implication is that components of the signal which do not meet these conditions (long spectral or temporal tails) will not be addressed by the initial conditions, and will therefore not be stored as  $\Omega$  is reduced, diminishing the efficiency. That said, control optimization goes a long way towards mitigating the required optical depths, and can be achieved through experimental iteration as well. In an experiment reported on just after Gorshkov et al. published their work on optimal control, an EIT based memory implemented in hot Rb vapor achieved 43% storage efficiency with an optical depth of only 24 [168], and the equivalence of optimizing the control shape for a given signal and of optimizing the signal shape for a given control was simultaneously demonstrated. Strangely, and contrary to the model of [117], the efficiency proceeded to saturate and then decline as the optical depth was increased further in this experiment, but the explanation for that phenomenon will have to wait another page or two. In trapped, cold atomic ensembles, high optical depths can be achieved using cigar shaped atom clouds also resulting in high efficiencies, for example 78% efficiency was demonstrated for an optical depth of 156 [169].

In the previous section I pointed out that the physical limitation to memory bandwidth is determined by how close the next nearest hyperfine transitions lie to the isolated  $\Lambda$ -scheme. This observation is true generally, however when the focus is broadened from rare-earth ion doped crystals to other  $\Lambda$ -systems like hyperfine-split D lines in alkali atoms this bound jumps from the order of megahertz to gigahertz. In fact, the bandwidth of EIT memories is rarely limited by the fundamental properties of the atoms. As can be read off of equation 2.3, the width of the transparency window depends on the available control power. Far enough into the Autler-Townes regime, they are approximately equal. Let us perform an estimate of how this condition translates into experimental parameters, as this is also relevant to the discussion of fidelity. The Rabi frequency of a transition is given by

$$\hbar\Omega = \mathbf{d}\mathbf{E}, \quad (2.4)$$

where  $\mathbf{d}$  is the dipole moment of the transition and  $\mathbf{E}$  is the electric field of the laser. Characterizing the laser beam experimentally, its power and beam shape can be measured yielding its intensity  $I$ . Then the field amplitude is  $E_0 = \sqrt{\frac{2I}{c\epsilon_0}}$ . In terms of

---

<sup>6</sup>Note that both the authors of [167] and [166] define their Rabi frequency as half of the experimentalists' value. The authors of [166] additionally halve their optical depths as is convention in theoretical descriptions. As noted previously, I have regularized these definitions to match measured values.

the order of magnitude, to obtain a Rabi frequency of 1 GHz on a Rb D<sub>1</sub> line transition in a beam with a diameter of 0.5 mm, between 100 mW and 1 W of laser power are required. This proves to be an awkwardly large amount of power, especially considering how small we have already assumed the spatial mode to be. Further, given the optical losses generally associated with fast switching or pulsing as is required to stop the light, a simple diode laser will require amplification to bring this Rabi frequency to bare on the atoms. These technical considerations are part of the reason why EIT memories often operate at lower bandwidths, on the scale of tens of megahertz, even though this does not reflect a physical limitation. The other challenge to employing such strong control beams is separating their light from stored single photons or weak signals. Continuing the order-of-magnitude estimate, a short control pulse designed to store a signal with a bandwidth between 100 MHz and 1 GHz will require a pulse energy around 1 nJ. The photon energy, again on the Rb D<sub>1</sub> line, is around 1.6 eV [115]. This means that such a control pulse contains about  $4 \times 10^9$  photons. This is an enormous amount of light to split from a single signal photon, even before taking into consideration any additional light-emitting atomic response to the control, like fluorescence or four-wave mixing. In the section of their review of quantum memories discussing this challenge, Heshami et al. [50] survey the typical range of the number of photons in read/write control pulses and give it very broadly as between  $10^6$  and  $10^{13}$ . This is the origin of the greatest hurdle to the fidelity of broadband quantum memories: readout noise. While only tangentially related to the memory protocol, after all they are sensibly designed and all yield unity fidelities under ideal conditions, no discussion of implementations, particularly in hot vapor, can sensibly omit this topic. Before moving on to describing further schemes I will therefore discuss noise and memory lifetime, particularly in hot vapors.

### **Readout Noise and Lifetime**

A significant portion of the experimental part of this thesis focuses on practical means of eliminating readout noise, so my main focus here is to illustrate the problem. Common to all schemes employing strong optical control pulses is the necessity of separating signal and control. Usually these are orthogonally polarized, so a polarizer is used as the first filter. The remaining control is typically filtered spectrally, exploiting the few gigahertz hyperfine splitting to remove the rest of the laser light. This combination of polarization and spectral filtering provides the required control rejection [170, 171, 172, 173], and around 140 dB suppression in total is not unusual. Admittedly, this sometimes requires unwieldy “stacks” of 7 etalons [172], but the problem is reliably solvable. Spatial filtering of the optical control is of course convenient in protocols where control and signal counterpropagate, but in co-propagating protocols even small


angles between signal and control have considerable drawbacks. For one, EIT shows strong angular dependence as a process on its own already [174], but moreover the resulting wave-vector mismatch  $\Delta\mathbf{k}$  causes a spatially varying phase of the stored spin wave in the medium. If the atoms move in between storage and retrieval, this leads to dephasing. For the  $j$ th atom the phase change is  $\Delta\phi_j = \Delta\mathbf{k}(\mathbf{r}_s - \mathbf{r}_r)$ , where  $\mathbf{r}_s$  ( $\mathbf{r}_r$ ) is the atom's position during storage (retrieval). This leads to a spin-wave-dephasing limited memory lifetime of [175, 176]

$$\frac{1}{\Delta\mathbf{k} \cdot \mathbf{v}} \sim \tau = \left( \frac{4\pi v}{\lambda} \sin \frac{\alpha}{2} \right)^{-1}, \quad (2.5)$$

where  $\lambda$  is the optical wavelength,  $v$  is the atomic velocity, and  $\alpha$  is the angle between the signal and control modes. For angles that produce significant filtering this is likely to be the limiting process to the overall memory lifetime in hot vapors. Dissimilarly in cold atoms, where the velocities are small, spatial filtering can be used without dephasing limiting the lifetime – like in [176] where an angle of  $13^\circ$  is used for this purpose. Decay of this kind is easily identified in data as it follows a Gaussian flank, in contrast to exponential state-decay or other forms of atom loss.

Despite the magnitudes involved, we have noted that filtering to remove the control from the read out signal after the memory is successfully employed rather routinely. The real remaining problem lies in noise arising from the atoms induced by the incident control. For processes that produce light which is spectrally separated from the signal we can again fall back on filters, but this spectral distinction cannot always be relied upon. The first potential source of unwanted signal-resonant photons traces its origin to residual atomic population in the storage state. In alkalis, with ground state hyperfine splittings  $\Delta\nu < 10$  GHz, we reach  $h\Delta\nu_{\text{hf}} \approx k_B T$  for temperatures in the range of 100 mK. It is therefore safe to say that at room temperature the states  $|g\rangle$  and  $|s\rangle$  will be equally populated initially. If nothing is done about this, the read out control pulse scatters from atoms in  $|s\rangle$  producing anti-Stokes photons at the signal frequency. Optical pumping must thus be used to empty  $|s\rangle$  before operating the memory. Moreover, this state preparation should last longer than the desired storage times, and, as is typical in vapor cells, if only some of the atoms available are addressed by the memory interaction it is desirable to pump a larger region so that atomic motion cannot effectively deteriorate the prepared state. Consider also that there are limits to how well optical pumping can work in vapor cells [177]. I will comment further on those considerations in section 5.4. This source of noise is negligible without pumping when  $h\Delta\nu_{\text{hf}} \gg k_B T$ , as is the case e.g. with memories implemented in diamond [178].

Even if all the atoms are initially in  $|g\rangle$ , the control can still couple them to  $|e\rangle$ , just off-resonantly by  $\Delta\nu_{\text{hf}}$  plus any additional detuning used, again producing fluorescence.



**Figure 2.6:** *left* A photon from the control laser, either in the write or read pulse, scatters inelastically to produce a Stokes photon, detuned from the signal by  $2\Delta\nu_{\text{hf}}$ , and an atomic excitation. *right* During readout another control photon scatters at the previously created atomic excitation, producing an anti-Stokes photon at the exact frequency of the signal.

In hot atoms with typical Doppler-broadening on the order of 500 MHz, the innate detuning  $\Delta\nu_{\text{hf}} \approx 10\Gamma$  leaves plenty of scattering opportunities for the number of control photons involved, and large additional detunings are required to suppress this coupling. This is the solution provided by Raman memories, which I will discuss imminently in the next subsection, but even for moderate detuning the situation is not quite as dire as it may initially seem. It turns out that, through a collisional relaxation process in the excited state, fluorescence noise is emitted at the natural frequency of the line even if the signal and control are detuned from it [179], but with the Doppler-broadened linewidth. Since the first observations and models in the '70s this appropriately named *collisional* fluorescence has also been studied in the context of memories [180], and has been shown to cause decoherence during the read and write operations [181]. That said, as the noise occurs at the natural frequency of the line, detuning does bring spectral filtering back on the menu. In cold alkali atoms this solves the problem rather neatly, as only a small detuning is required to escape the natural excited state linewidths of  $< 10$  MHz. In hot atoms the situation is not improved as rapidly, but it does mean that there is some wiggle room with how much detuning really is required to diminish the role of fluorescence.

There is a third important source of control induced noise: four-wave mixing (FWM). This two step process, illustrated in figure 2.6, produces additional undesirable photons at the signal frequency. In reviews on quantum memory, e.g. [50], it is common to find it pointed out that if  $\Delta < \Delta\nu_{\text{hf}}$ , then the memory interaction experiences stronger coupling than FWM which allows it to dominate over the noise. The corollary that FWM only matters in the case of far-off-resonant Raman memories with  $\Delta \gg \Delta\nu_{\text{hf}}$  is false. This is due to the optical depths required for efficient storage. At high OD the

FWM process can dominate EIT, even at  $\Delta = 0$  [182, 183]. In the context of memories, this is best understood as interference producing gain or loss<sup>7</sup> in the signal field, and from their combined EIT-FWM memory model [184] estimate the trouble to begin at optical depths around 25. Indeed, it is with this model that an explanation is found for the saturation and subsequent decline in memory efficiency with optical depth these same authors previously observed [168]. The functional limit to the optical depth before FWM noise grows exponentially large was found in [185] to be  $\text{OD} > \frac{2\Delta|\Omega'|}{\Gamma|\Omega|}$ , where  $\Omega$  is the Rabi frequency on  $|s\rangle \rightarrow |e\rangle$ ,  $\Omega'$  is the Rabi frequency on  $|g\rangle \rightarrow |e'\rangle$ , and  $\gamma_{ge}$  is the natural linewidth. To clarify the notation, in the simple case of one allowed transition  $|e'\rangle = |e\rangle$ , but if there are multiple available excited states these need to be accounted for both here and in defining an effective  $\Delta$ . Concretely it is also shown that FWM gain is harmful; it necessarily reduces the fidelity of EIT memories unless it is already below  $\frac{1}{\sqrt{2}}$ , in which case the memory is unsuitable for quantum storage right from the start.

It may now seem that this discussion of noise did not yield a lot of information that we can directly translate into the language of our figures of merit. This is in part due to the practical fact that noise forces trade-offs in these figures. Less control power may improve a memory's noise performance, at the cost of its bandwidth or its efficiency. Nevertheless, there is one simple estimate we can make for the effect of incoherent readout noise on the quality of a stored single photon. Referring back to equations 1.10 and 1.11 we can write,

$$g_{\text{retrieved}}^{(2)} \approx \frac{g_{\text{input}}^{(2)} \times \text{SNR} + g_{\text{noise}}^{(2)}}{\text{SNR} + 1}, \quad (2.6)$$

which is just an average value of the autocorrelations of the input and the noise. Due to any multi-photon component of the input skewing the detection on non-number-resolving detectors in Hanbury Brown and Twiss configuration, this estimate is a bit sloppy. In a proper calculation where the fields from disparate sources are combined on an imagined beam splitter, following arguments in [186], the resulting relative multiphoton emission probability in the presence of incoherent noise is,

$$g_{\text{ret}}^{(2)} = \frac{(N_{\text{ret}} - N_{\text{noise}})^2 g_{\text{input}}^{(2)} + 2N_{\text{noise}}(N_{\text{ret}} - N_{\text{noise}}) + N_{\text{noise}}^2 g_{\text{noise}}^{(2)}}{N_{\text{ret}}^2}. \quad (2.7)$$

The difference in these formulas vanishes when  $g_{\text{input}}^{(2)} \rightarrow 0$  and  $g_{\text{noise}}^{(2)} \rightarrow 2$ . Note that

---

<sup>7</sup>Gain is of course a particularly insidious phenomenon, as without an analysis of the fidelity or photon quality, an unwitting fan of quantum storage could easily convince themselves that their memory functions well even though it just produces noise. I believe that it is this, and not the popular sparing of words in physics, that have coined the effect FWM gain.

this equation does not model the effect of all the noise sources discussed so far, in particular FWM gain requires a coherent model for the autocorrelation of the memory output [187]. Nevertheless, this gives us at least some idea of the problems noise poses to the project of storing photons, so let us now return to the final figure of merit on our list.

Achievable storage times depend on the physical memory implementation rather than just the scheme. It comes as no surprise that for EIT based memories doped solids excel in this regard just as well as they do in the photon echo case, after all they do not experience significant atomic motion. With high population lifetimes, for instance around 100 s in  $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ , the challenge becomes to apply dynamic decoupling techniques (e.g. ZEFOZ, see above) optimally so as not to suffer from decoherence due to stochastic magnetic interactions with the host material. This is challenging as the applied magnetic field can split the spectrum into over 1000 separate transitions, but the development of evolutionary algorithms to solve the optimization problem with experimental feedback has led to storage times almost matching the population lifetimes [188]. In cold atoms optical lattices can be used to prevent atomic motion as well, and analogous magic field and microwave dynamic decoupling protocols have been used to achieve similar minute-time scale lifetimes [189]. What about in hot alkali vapor where atomic motion is a given? The ballistic atomic velocity in thermal equilibrium for atoms of mass  $m$  at an absolute temperature  $T$  is  $v_{\text{ball}} = \sqrt{k_{\text{B}}T/m}$ . Depending on the time scale of interest and mean free path in the vapor, the mean atomic velocity  $\langle v \rangle_T = \sqrt{8k_{\text{B}}T/\pi m}$  may be a better estimate due to collisions. Either way, this yields a mean time-of-flight of an atom through a laser beam with radius  $r$  of  $\tau = r/v$ . After this time the atom is either no longer addressed by the read pulse, or, if the control beam and vapor cell have the same diameter, it collides with the cell wall which randomizes its state. To combat this the atomic velocity can be slowed, that is made diffusive, by the use of a buffer gas. The atom-laser interaction time then becomes [166]

$$\tau^{-1} = 2.405^2 \frac{D}{r^2 + 6.8\lambda_{\text{mfp}}r}, \quad (2.8)$$

where  $D$  is the diffusion constant for the alkali through the buffer gas and  $\lambda_{\text{mfp}} = 3D/\langle v \rangle_{T, \text{alk-bg}}$  is the mean free path, with the mean relative velocity now defined as  $\langle v \rangle_T = \sqrt{8k_{\text{B}}T/\pi M}$ , where  $M$  is the reduced mass of the colliding system. But even if an atom moves diffusively, sooner or later it will encounter the cell wall. To mitigate the severity of the wall collision, an anti-relaxation coating can be applied. Standard materials, meaning that they are easily commercially available, include paraffin, with which atoms can tolerate around  $10^4$  wall collisions before becoming depolarized, and OTS (octadecyltrichlorosilane), which is about an order of magnitude worse in terms of the number of collisions but survives higher cell temperatures [190]. The state of the




art consists of alkene chains with which minute-long transverse spin-relaxation times have been demonstrated [191, 192, 193].

Collisions between alkalis can also cause decoherence through the process of pairwise spin-exchange. Therein the electron spins of the colliding alkalis rotate with respect to their (conserved) combined spin. The time between these collisions takes the form  $\tau = (\sigma_{se}n\langle v \rangle_T)^{-1}$  where  $n$  is the number density of the vapor and  $\sigma_{se}$  is the spin-exchange cross section. Typical values of  $\sigma_{se}$  lie around  $1 \times 10^{-14} \text{ cm}^2$  to  $3 \times 10^{-14} \text{ cm}^2$  across the alkalis [194]. The rate of decoherence due to these collisions depends on the Zeeman level populations, but with one notable exception the order of magnitude remains the same [195]. This decoherence can be suppressed in numerous ways, for example by ensuring that the rate of spin-exchange exceeds the Larmor precession frequency, an operating regime called spin-exchange relaxation-free (SERF) famous for its use in magnetometry [196]. Alternatively or additionally, it is possible to work with the aforementioned notable exception: the stretched state of the atoms,  $|m_F| = I + \frac{1}{2}$ . Here spin exchange is forbidden due to conservation of angular momentum. Recently in 2018, a memory implemented in hot Cs vapor utilizing both wall coatings and spin-exchange suppression mechanisms achieved a storage time of 430 ms [197], demonstrating that these measures to combat decoherence work practically and that long storage times are within the reach of hot vapor memories broadly if the effort is made to implement them. Thus this interlude on general considerations applicable to memories in hot vapor vis-à-vis these figures of merit concludes optimistically, albeit without having reduced noise concerns to a negligible level. Let us therefore turn to a few further quantum memory implementations that tackle this issue at the level of protocol design.

## Raman

Raman scattering is an inelastic two-photon process. In the just discussed context of four-wave mixing (see figure 2.6) we have encountered its “normal” form, albeit with reversed roles of the states of the lambda system. To see how this interaction can be used for storage, let us first return to our regular notation with figure 2.7. A weak signal or photon  $\mathcal{E}$ , detuned from the  $|g\rangle \rightarrow |e\rangle$  transition by  $\Delta$  is incident on a lambda-system. A strong control  $\Omega$  tuned to two-photon resonance, meaning that it is also detuned by  $\Delta$  but from its own transition  $|s\rangle \rightarrow |e\rangle$ , is simultaneously applied. The control produces the stimulated emission of a Stokes photon and thus maps the signal into the usual ground state spin wave, with one undetermined atom transferred from  $|g\rangle$  to  $|s\rangle$ . The detuning for this process can be chosen to be large compared to the photon bandwidth and the excited state linewidth, ensuring that  $|e\rangle$  is never occupied. This coupling can then be understood as the absorption of the signal by a



**Figure 2.7:** *left* In a two-photon resonance process a signal photon  $\mathcal{E}$  and control pulse create an atomic excitation by the stimulated emission of a Stokes photon into the control beam. *right* For readout the control is applied again and the reverse process reproduces the signal photon.

virtual excited state created by the control. This picture is essentially the same as what we derived using dressed states in the EIT case, identifying the virtual state of the Raman transition with  $|-\rangle$ . The distinction between the cases then lies in whether the signal is tuned in between the dressed states (EIT) or into resonance with just the one of them consisting mostly of  $|s\rangle$ . I owe this concise delineation of the processes to the explanation Josh Nunn gives in his thesis theoretically treating photon storage in Raman memory [198], and have adopted it as I feel it gives more insight into the transition between regimes than the mere observation that EIT occurs at and near resonance and that Raman storage is the far detuned limit. In particular, the physical difference is that Raman storage is not an interference effect of absorption pathways, and that there is no reduction of the group velocity of the signal [55].

The theoretical description of Raman storage began with a scheme to transfer quadrature squeezing between light and atoms [199], and was extended to the problem of storing single photons later on [200, 198]. The first demonstration with weak light pulses in hot Cs vapor [60] emphasized the memory bandwidth, as gigahertz broad signals were successfully stored right away. A follow up experiment published in 2011 [201] examined the memory’s unconditional noise floor, measured to be around 0.25 photons, and observed that indeed it was not plagued by the collisional fluorescence noise that had, at the time, raised concerns about the viability of hot vapor EIT memories (see [180, 181] and the previous section). Despite observing considerable four-wave mixing noise, the authors of [201] were optimistic in their discussion that the problem could be solved. Four years later however, when the scheme was tested with single photons, this noise source dominated the photon statistics of the readout [187], and it became evident that its suppression was essential. A method to do so placing the memory in a low-finesse ring cavity resonant with the memory interaction (signal and

control) and anti-resonant with the noise process was developed [110] and implemented [172], but came at the price of considerable technological complexity due to the triple resonance condition and stability requirements, for example temperature stability to within 0.1 °C (see supplementary material to [172]). Nevertheless, these efforts were rewarded with a solid order of magnitude improvement in the noise performance. Simultaneously, other methods to suppress this noise source were being proposed. Of those applicable both to Raman and EIT memories, a particularly interesting solution due to its simplicity was to just have the noise be absorbed by another populated transition in the atoms [202]. Originally this was proposed to be a transition in a second isotope present in a natural vapor cell, but an implementation published in 2019 showed that the same noise absorbing effect could be achieved by operating the memory at a detuning equal to twice the hyperfine splitting [203]. In summary, the project of eliminating noise in hot vapor lambda memories has spawned many ideas, and much progress has been made. The discussion of other figures of merit in hot vapors from the previous section naturally applies to hot vapor Raman memories as well. Furthermore, Raman memories have seen implementations in many other systems, including, for a somewhat exotic example, the rotational states of hydrogen molecules. This implementation has been shown to work with a signal bandwidth of 7 nm (about 3 THz at the operating wavelength) and uses polarization selection rules to avoid FWM noise [204]. Simultaneously the memory lifetime is around 80 ps, which is merely to imply that its use is highly local, perhaps in frequency or bandwidth conversion [205]. Rather than delving further into the cornucopia of possible implementation examples, I will now move on to yet another kind of memory scheme addressing noise problems in hot vapor.

### 2.1.5 Ladder Memories

Recall that at the beginning of the section on lambda memories we set out to gain coherent control over a hypothetical two-level memory by adding a third state to the system. The choice of a metastable ground state was motivated by such a state's lack of radiative decay, and in the fullness of time we found that indeed memories making this choice could achieve long lifetimes. Nevertheless, with regards to our initial objective, we could also have chosen another excited state in a higher orbital. This type of three-level system is called a cascade or ladder scheme, and a small number of memories have been implemented in such a scheme [206, 207]. In particular, [206] uses the cesium D<sub>2</sub> line for its  $|g\rangle$  and  $|e\rangle$  states, and the state 6D<sub>5/2</sub> as  $|s\rangle$ , while [207] does similarly with the rubidium D<sub>2</sub> line and the 5D<sub>5/2</sub> state. The radiative lifetimes of these states are 400 ns and 260 ns respectively, and measured lifetimes turned out to be shorter still due to the wavelength difference between signal and control leading

to spin wave dephasing, akin what we saw in equation 2.5 due to misalignment of the signal and control modes. Recently, a quite general technique to protect collective excitations from inhomogeneous dephasing by dressing them with an auxiliary state was developed [208], and demonstrated therein using the ladder memory from [207]. But even if this technique is used, the short radiative lifetimes certainly will be limiting – so what do we gain? Well, in exchange it becomes trivial to remove the control laser from the signal with low absorption color filters, and many of the intrinsic noise sources we previously identified, fluorescence, four-wave mixing, and residual occupation of the storage state, are effectively eliminated. Collisional fluorescence during read-in or read-out might still occur. This possibility is addressed in both [206], who note that a detuning can be used, and [207], who state that they minimized it by leaving plenty of time after pumping for the excited state to relax. Although both are operated with detuning, no narrowband spectral filtering is used in either implementation to achieve the low noise. Moreover, the physical limit to the bandwidth in this scheme is now not given by how near another hyperfine state lies to  $|s\rangle$ , but by how near another orbital state lies. In practice this means the limit will be presumably given by the control, either due to its bandwidth or available intensity, but the physical limit to the acceptance bandwidth is increased by many orders of magnitude over lambda storage. These advantages have given rise to a surprising application for hot vapors – the noise performance is so good in ladder memories that storage therein, application in mind now called a noise-free buffer by the authors, could improve the indistinguishability of photons from a quantum dot at a lower brightness cost than conventional filtering schemes [209].

### 2.1.6 Beyond Three-Level Systems and Nuclear Coherence

In terms of three-level atoms we have now almost exhausted the possible configurations. The inverse of the  $\Lambda$ -system, the V-system, remains, but this type does not seem to offer notable advantages. Instead it inherits the lifetime issues of ladders and most of the noise issues of lambda systems. In four-level N-systems, with two applied control fields, a memory interaction based on electromagnetically induced absorption (EIA) begins occurring which was only recently studied in greater detail [210]. Five level M-systems and beyond are home to all kinds of laser induced coherence effects including EIT, EIA, and coherent population trapping, but also to higher orders thereof analogously named EITA, EITAT, and so on [211]. These do not seem to have been studied in the context of storage explicitly. Rather than speculating, I want to briefly mention a particularly interesting “extension” to light storage in three-level atoms, the transfer of a stored electronic spin coherence to a nuclear spin coherence.

Nuclear spin coherence times tend to be significantly longer than electronic ones.

The electron shells shield the nuclei from the environment, and their response to magnetic fields is a factor  $10^3$  weaker. We briefly encountered this fact with the 6 h lifetime record [139] achieved in a CRIB memory with optically addressable nuclear spins. In mesoscopic systems as well, quantum dots in particular, reversible mapping between electronic and nuclear coherence via spin-exchange has been proposed as a method for creating a long lived memory [212], but here the initial electronic decoherence rates are faster than in alkalis and consequently the contextual notion of “long-lived” still only implies timescales of seconds. These, and quite generally other nuclear coherence times, particularly at room temperature, are dwarfed by those achievable with the exceptionally well shielded nucleus of  $^3\text{He}$ . In  $^3\text{He}$  nuclear relaxation times of thousands of hours and coherence times of 60 h have not just been observed, but due to its wide spread scientific applications large quantities of highly polarized gas are routinely produced, for example for magnetic resonance imaging of lungs [213]. In 2005 it was first proposed to use the same mechanism commonly used to polarize helium in the first place, metastability exchange (ME) collisions, to transfer a squeezed state of light via the optically accessible metastable triplet state  $2^3\text{S}_1$ , populated by electric discharge, to the helium ground state [214]. Moreover, a scheme has been proposed to map a Faraday interaction between light and metastable helium states to the ground state via ME collisions [215]. This has direct relevance for quantum memory via a continuous variable scheme [216, 217], a topic I have not discussed in detail previously<sup>8</sup> but has seen successful implementation [219]. Finally, an interfacing mechanism based on weak spin exchange collisions between alkali vapors and helium has been described [220] and implemented [221], with an exchange time scale of a few milliseconds. The implication is that in a vapor cell filled with an alkali and helium, first regular lambda storage of light could be performed in the alkalis without requiring extremely long lifetimes, after which the collisional interface with helium preserves the stored state for hours in the helium ground state [222].

## 2.2 Review of Single Photon Sources

To open the discussion of single photon sources, we must first establish what a photon is. Following Dirac, photons are theoretically conceived of as the elementary excitations of

---

<sup>8</sup>Implicitly, my exposition of quantum memories has considered the storage of discrete variables. Information is encoded in where the photon is, e.g. spatial mode 1 or 2 (path encoding), its polarization, which is cast as one of two orthogonal basis states, in its arrival time, which is discretized into late and early time-bins, or perhaps in opposite orbital angular momenta, which are also discrete [218]. However, as noted in chapter 1, the effective aspect of the quantum lies in stronger-than-classically-allowed correlations, which leaves the door open for such correlations, squeezing, in continuous variables such as phase or amplitudes quadratures. An analogy to the distinction between digital and analogue memory lies close at hand, but ultimately escapes the scope of this review.

a quantized electromagnetic field [223]. This formalization of the notion of a single light particle defines, via the choice of the quantized field mode  $k$  with frequency  $\nu_k$ , that the photon is a monochromatic excitation with the energy  $E = h\nu_k$ . In practice, the object of study usually seems somewhat different. It is the result of some physical process understood to produce photons one at a time, but they are localized in spacetime and thus not truly monochromatic. The solution to this apparent discrepancy is offered by considering the coherence of the electromagnetic field, first-order coherence signifying the occupation of a single mode even without it being monochromatic [224]. This approach naturally allows for localized photon wave functions, wave-packets, and a theory of photon wave mechanics. It should be noted however, that these are simply different formalisms for the same quantum field theory, and that conversions are always possible [225]. Considering higher-order field coherence produces a concrete set of observable properties, the number statistics measured in correlation experiments, by which the nature of a field can be categorized into groups like thermal light, coherent light, and single photons [226]. Moreover, it is in these correlations that the quantum usefulness of photons lies. Therefore, a compromise widely adopted is an operational definition founded on these observable and essential properties: single-photon states are those that produce the photon-number statistics of single photons when measured with photon detectors [227]. Then more generally, a photon-number state is one with an integer number of photons.

We can now imagine an ideal source of such single-photon states. It would emit indistinguishable photons, one at a time and on-demand, with certainty ( $p = 1$ ), no chance of additional emissions, and at a rate bounded only by the duration of the single photons themselves [228]. Moreover, the condition of indistinguishability should apply to the photons emitted by physical copies of the source, i.e. it must be preserved when scaling the number of sources [218]. For multi-mode or multiplexed sources emitting many distinct photon modes, these conditions apply to each of the modes individually. Given the ambitious assumption that the modes are practically separable, an ideal multi-mode photon source is the same thing as many single mode sources. From this defining idealization we can read off most of the appropriate figures of merit for single photon sources. I will do this now, commenting on their significance in the context of storing the single photons in memories where applicable. Then I will briefly review three kinds of photon sources: atomic sources in the form of read-only memories, mainly due to this platform's similarity to memories, as well as quantum dots and spontaneous parametric down-conversion (SPDC) sources in non-linear crystals, both due to the attractiveness of interfacing these sources with our kind of atomic memory. More comprehensive overviews of platforms are summarized in table 1 of [227] and table 2 of [218]. The requirements various quantum applications set for these figures of merit are collected in table 2 of [229].

### 2.2.1 Figures of Merit

Single photons from some sources, for example SPDC or FWM, have spontaneous processes at their origins. We can divide useful single photon sources into two types, those that (in principle) generate photons deterministically upon a trigger, usually isolated systems of single emitters, and those that generate photons probabilistically, but in pairs. Naturally, only the former can meet the ideal captured in the term “on-demand” in our definition, however if a spontaneously generated photon pair can be split, then one photon can be detected to indicate the existence of the other. By one convention,<sup>9</sup> the photon that is detected is labeled idler while the other is called signal. This process is called *heralding*. The technical implications of using such a source are significant. Effectively, whatever apparatus the photon is to be interfaced with must now be able to respond to a trigger from the photon source. The source becomes the experimental master clock. Whether this requirement is acceptable at all depends entirely on what the photons are to be used for. The distinction is therefore less a figure of merit than a binary choice, either a heralded source can be used or it can not. Simultaneously, it is important not to take the moniker of a deterministic source too seriously, as it is not to imply that the probability of obtaining a photon from the source upon a trigger is 1. The distinction this label is supposed to imply is a different one, namely that the experimenter either is in control of when the photon is emitted by the source (on-demand) or at minimum knows the time of emission in advance, if a photon is indeed emitted. The latter can also apply to heralded sources when they are pumped with sufficiently short pulses. For this reason it makes sense to distinguish the figures of merit between these kinds of implementations. Note also that a spontaneous photon pair source interfaced with an on-demand memory for photons could be recast as a deterministic photon source.

#### Indistinguishability

Indistinguishability is the photon source counterpart to the memory fidelity, and as in section 1.2.2 it can be defined as the quantum state overlap between photons produced at different times or in different sources. For characterizing one or two sources, the measure employed for this figure of merit is the Hong-Ou-Mandel visibility we encountered in equation 1.13. The HOM effect can also be generalized to beam splitters with  $n$  ports for tests of the indistinguishability of  $n$  photons [232]. On the scale of multi-mode networks, so called Sylvester interferometers [233] have been shown to enable near optimal indistinguishability tests [234]. Indistinguishability is of

---

<sup>9</sup>This convention is motivated as *signal* seems to be a more appropriate term for the photon we want to do something with, and it is also the one followed by Boyd [230]. A common alternative is to call the photon with the higher frequency the signal as in Powers and Haus [231].

fundamental importance to ensure that quantum interference between these photons occurs as expected by protocols extracting an advantage from entanglement, for instance to enable the violation of a Bell inequality [235]. Practically, the operation of a storage interface over time might also explicitly assume that every photon to be stored is the same, for instance in setting a fixed frequency, duration, and intensity of the control laser.

### State Accuracy and $g^{(2)}(0)$

Most generally, the output of a photon source can be characterized by a full quantum state tomography [236] and the result can be compared to the single photon Fock state. Fortunately, there are also simpler methods to check whether a source accurately produces single photons. The analysis of light by its coherence properties [237] and the identification of arbitrary states of radiation by their photon statistics [238] was famously performed by Roy Glauber in 1963. The textbook by Christopher Gerry and Peter Knight [65] presents most of the relevant elementary results, so I am once again refraining from regurgitating everything here and will only present what I consider essential to the later chapters. Crucially, the question arises whether a machine dubbed single-photon source actually outputs single photons by our statistical definition. Upon detecting a photon at time  $t$ , as a function of the delay  $\tau$  until the next photon detection, the second-order correlation function of the radiation field is

$$g^{(2)}(\tau) = \frac{\langle \hat{E}^{(-)}(t) \hat{E}^{(-)}(t + \tau) \hat{E}^{(+)}(t + \tau) \hat{E}^{(+)}(t) \rangle}{\langle \hat{E}^{(-)}(t) \hat{E}^{(+)}(t) \rangle \langle \hat{E}^{(-)}(t + \tau) \hat{E}^{(+)}(t + \tau) \rangle}. \quad (2.9)$$

Here  $\hat{E}^{(+)}$  and  $\hat{E}^{(-)}$  are the positive (absorbing, proportional to the photon annihilation operator) and negative frequency components of the electromagnetic field, respectively. For a single mode field, this reduces to

$$g^{(2)}(0) = \frac{\langle \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} \rangle}{\langle \hat{a}^\dagger \hat{a} \rangle^2} \quad (2.10)$$

$$= 1 + \frac{\langle (\Delta \hat{n})^2 \rangle - \langle \hat{n} \rangle}{\langle \hat{n} \rangle^2}, \quad (2.11)$$

as in equation 1.11 which is formulated in terms of detection probabilities. This parameter yields two distinct possible indications of non-classical light, one being sub-Poissonian statistics indicated by  $g^{(2)}(0) < 1$ , and the other being photon antibunching indicated by  $g^{(2)}(0) < g^{(2)}(\tau)$ . With these measures it can be experimentally quantified, independently of losses, what the multi-photon component of the state is. Obtaining  $g^{(2)}(0) = 0$  in a Hanbury Brown and Twiss experiment unambiguously indicates a light field with no more than one photon. Mainly in the domain of solid-state sources, a



popular name for this measure is single-photon purity (see e.g. [229, 239]). As what we are describing is intended to be a source of single photons, and as the word pure already does a lot of heavy lifting in quantum mechanics, I prefer the notion of single-photon *state accuracy* to mean what is also called the “single-photon nature” [227] (both terms *and* the scare quotes are used in this source). Concretely,  $g^{(2)}(0)$  expressing the relative probability of multiple photon emission to single photon emission has the obvious interpretation as an inaccuracy in producing single photons. Like purity, state accuracy thus refers to the measure approaching zero, but that there is one target is more explicitly acknowledged even when the prefacing *single-photon* is omitted. In the case of heralded sources, the relevant correlations involve three photons. Only the correlation function between signals conditioned on the detection of an idler,  $g_c^{(2)}$ , can go to zero. In general, the details depend significantly on the heralded source implementation. To get a taste of the specificity involved, consider that a complete description of the photon correlations in a continuously pumped and cavity enhanced SPDC source was published only recently in 2020 [240].

### Efficiency

For deterministic sources, the efficiency can be conceptually divided into an excitation probability and an extraction or collection probability. Some procedure, for example an optical  $\pi$  pulse, is applied to the photon source, whereupon it should output a photon. The excitation efficiency corresponds to how likely this procedure is to generate a photon within the source. Before the generated photon can be used it has to come out of the source. Excluding the special case of integrated optics on chips or in waveguides, a reasonable standard for general usability is to have the photon come out of a single-mode fiber. This is typically necessary for detection anyway, and implies that the photon can be interfaced with anything that accepts a fiber input. The probability of the photon coming out of this collection fiber upon successful source excitation is the extraction efficiency. It is common to find this latter efficiency defined at an earlier element, for example after the first lens [239]. This delivers no guarantee that only a single spatial mode is considered, and leaves open the question of how these photons now get to where they are needed, so I reject this approach. Together these efficiencies yield an end-to-end efficiency, the success probability of the idealized “single-photon button”.

For heralded sources the extraction efficiency requires some nuance. A single “arm” of spontaneously generated single photon pairs exhibits thermal statistics [65]. The relevant measure to obtaining a single photon is therefore how well the presence of the signal photon is heralded by the detection of the idler photon. Notice the asymmetry this implies – collection inefficiencies or losses in the idler arm negatively affect rates,

but not the heralding efficiency. To optimize it, only the signal arm must be collected as effectively as possible. Enforcing the same collection and usability standard to both kinds of sources, the heralding efficiency is thus defined as the probability of having a signal photon come out of a single-mode fiber leaving the source upon the detection of the idler photon. The extraction efficiency of the idler in a heralded source can have a small secondary effect on the state accuracy, determining how much uncorrelated background is present in the signal arm relatively. After all, the loss of the idler to condition the signal on leaves its statistics thermal.

For memories with readout noise the efficiency of a source is particularly important. Every time the memory protocol is performed without a photon present pure readout noise is added to the memory output. The bound which needs to be set depends on the final application, and is minimally determined by the inverse signal-to-noise ratio of the memory for an ideal input.

### **Optical Compatibility, Modes**

The wavelengths and bandwidths of photon sources need to match the acceptance parameters of what they are interfaced with. This observation seems somewhat trivial considering that it also applies to, say, the dielectric coatings making mirrors reflective. Of course what I am really thinking of concretely are memories implemented on atomic and ionic lines. One possible option to ensure compatibility is to implement a photon generation scheme using the same optical line the memory operates on. Solid state sources of photons can, however, offer significant advantages in efficiency, state accuracy, and photon rates. It thus becomes a fabrication challenge to match the emission wavelength of the source materials to atomic lines, and an optical engineering challenge to match the bandwidth – either directly, building the source in a cavity, or through filtering the emissions. Frequency and bandwidth conversion schemes add to the possible options.

Before specific measures to address this issue, photon sources can also be multi-mode in a variety of ways. This is fundamentally undesirable as the modes make the photons distinguishable. If multiple modes are to be useful they must be separable physically and characterized by the figures of merit individually, i.e. one single-photon machine reliably producing two separable modes of single photons is by our definition actually two single-photon sources. I have already introduced single-mode fibers into the efficiency definition enforcing a single spatial mode, as required by our conceptual ideal. If the photons emitted by the source naturally are spatially multi-mode, then this fiber will act as a mode filter, regularizing the efficiency to the desirable mode and ensuring that our figures of merits measure the same property across implementations. Spectrally multi-mode sources require the same treatment, for example using filtering

etalons. Beyond a direct measurement with a sufficiently high resolution spectrometer, an attest to the spectral mode purity of single-photon sources is whether the time-bandwidth product of the photons is Fourier transform limited [241]. In heralded sources a measurement of the unconditional (not-heralded) autocorrelation also yields information about the mode number, the expected value in the low gain limit (see section 2.2.4) goes as [242]

$$g_{x,x}^{(2)} \approx 1 + \frac{1}{N}, \quad (2.12)$$

where  $N$  is the number of modes and a result of 2 (thermal statistics) is expected for a single mode.

### Rates, Brightness

To quantify how many photons come out of a single photon source over time, various definitions of rates and brightness are used. These vary across communities, and may be given with normalization to the spectral width, pumping power in heralded sources, or both. This can be useful for sources that are tunable over an operating range with known and characterized relations between different figures of merit within that range, for example and to first order, the pumping (laser) power in pair sources linearly scales the rate of pair generation and  $g_c^{(2)}$ . Deterministic sources have clock cycles given by the trigger or pulse rate. In this case brightness can be defined as the probability of having a photon at the source output per clock cycle. Obscuring the clock cycle in this measure is less than ideal in the context of interfaces, as this value needs to be matched between systems. For spontaneous sources there is no external clock to make reference to, so in that case the most useful metric is the rate of detected idler photons per second, which directly corresponds to the required clock cycle of any downstream interfaces from the source.

### Multiplexing Capability

Heralded photon sources are subject to both practical and fundamental [243] limits to their single-photon state creation probability, and generally speaking no single, real source currently fits the ideal envisioned. At some point we have to wonder whether the remaining insufficiencies really will be overcome with further improvements to single devices, for there is an alternative. Real photon sources' approximations of ideal deterministic sources can be improved by multiplexing [244]. Approaches include spatial multiplexing by building more sources or collecting different spatial modes from one device, temporal multiplexing by storing the photons, and spectral multiplexing by spectrally resolved detection of the idler. Switches are then used to combine the multiplexed source's output in a single mode. Due to the plethora of possible examples

and necessary assumptions, I direct attention to table 1 for improvements achievable through multiplexing, and table 3 for implementations of source multiplexing, in a recent topical review [244].

### 2.2.2 Atomic Sources and Read-Only Memories

There are a multitude of processes that produce single photons in atoms, so they have always been an obvious choice for a source of them. The earliest source of photons with which experimental distinctions between classical and quantum field theories were conclusively tested was based on an atomic cascade [245]. This test used  $^{202}\text{Hg}$  atoms, excited by electron bombardment, which decay on the cascaded system  $9^1\text{P}_1 \rightarrow 7^3\text{S}_1 \rightarrow 6^3\text{P}_1$  to produce one green and one blue photon. Cascades in other mercury isotopes, as well as one in calcium excited with krypton-ion and Rhodamine 6G dye lasers, were used in early tests of Bell's theorem [246], including the famously successful Bell tests by Alain Aspect et al. in the early '80s [247, 248]. Sources based on single trapped emitters saw development along with laser cooling and cavity QED experiments [249], and are still achieving some of the better state accuracies around [250], albeit at comparatively low rates and efficiencies. Modern atomic pair sources also include ones based on spontaneous four-wave mixing in ensembles [251, 252]. These sources make the best atomic analog to spontaneous parametric down-conversion sources in non-linear crystals, and when the latter is performed in cavity for compatibility to atomic lines the efficiencies and rates become quite comparable, but without either achieving the state accuracy of isolated single emitters. Moreover, atomic sources are usually compatible with other systems composed of the same kind of atoms without additional work. Finally, as atoms of a kind are identical, they are also at least capable of emitting identical photons, if the generation process is sufficiently noise free.

The most famous modern atomic photon source, and a fairly unique one at that, is that which is now commonly called the read-only memory [253]. It features in the Duan-Lukin-Cirac-Zoller (DLCZ) scheme for quantum repeaters and long distance quantum communication [254], and has thus become the staple example. This kind of source is heralded, but with a variable delay of the photon emission achieved as follows: The write process occurs through Raman scattering of a laser pulse in a lambda-scheme, which is described earlier in this chapter. After a first scattering event, a Stokes photon heralds the creation of an atomic excitation. The source, acting just like a memory except for the fact that the ground state spin wave was prepared spontaneously through scattering as opposed to deterministically read-in from a signal photon, is now ready. A read pulse can convert the spin wave into a photon on-demand within its lifetime. Implementations quickly followed the DLCZ proposal [255], and, as the archetypal system, they were also at the heart of a demonstration of a fully

heterogeneous quantum interconnect a few years ago [256].

### 2.2.3 Semiconductor Quantum Dots

Semiconductor quantum dots are nanoscopic structures consisting of one semiconducting material with a small bandgap, embedded within a different host material with a larger bandgap. The most widely investigated optical QDs at time of writing are made by depositing InAs or InGaAs on GaAs. The lattice mismatch causes strain, leading to the formation of small islands in a process called Stranski-Krastanov growth [257], or alternatively “self-assembly” [258]. These are then capped with more GaAs producing highly localized potentials for charge carriers due to the difference in the materials’ bandgaps. Even a simple model of a finite, 1D square potential, supplied with realistic length scales (10s of nanometers) and the effective electron mass of the material, can accurately predict the existence of a small number of bound states in these dots. These can be addressed with optical dipole transitions promoting an electron from a valence level into a conduction level and leaving a hole in its place. The result resembles a two level atom with a ground state consisting of full valence and empty conduction levels and an excited state of one such electron-hole pair, which is called an *exciton*. The described layers are often embedded in a semiconductor heterostructure, which may be gated for electrical control – for example to load a single electron into a conduction level before optical excitation producing a charged exciton, or to shift the energy levels by the Stark effect. The resonant wavelength of these quantum dots can be controlled in the range 900 nm–1200 nm by annealing after their growth [258]. Other materials and growth methods exist. Of particular interest to alkali interfacing projects are dots made by etching nanoscale holes in AlGaAs and letting GaAs diffuse into them before capping, as these dots emit around the rubidium D lines [259]. As fabrication techniques for these GaAs/AlGaAs, from here on “Rb-like”, dots are newer, and the dots themselves are not quite as commonly investigated as the longer wavelength ones, not all techniques for maximizing the figures of merit developed in “conventional” dots have been implemented in Rb-like ones. Interfaces with atoms are of general interest, and delaying light is a classic experiment demonstrated in Rb [260] but also Cs interfaces with conventional dots [261, 262]. Nevertheless the systems are sufficiently different that it is worth distinguishing between them until we know how their performances stack up.

Typical radiative lifetimes of quantum dots are around 1 ns, and their bare count rates have an edge over other solid-state single-photon emitters with similar lifetimes comparatively reviewed in [229]. This is despite the fact that a recurring difficulty is posed by the high refractive index of GaAs. It leads to a probability of only about 2% for photons to exit the semiconductor from one end facet [239]. To improve upon

this, the dots' environment can be engineered by embedding them in an optical cavity or waveguide. Table 1 of [239] collects fairly recent state-of-the-art performances of quantum dots using many different kinds of photonic structures to enhance the collection efficiency. Dots in microcavities have achieved particularly impressive end-to-end efficiencies including fiber-coupling, which is the most stringent metric among the ones commonly reported for the efficacy of light collection from quantum dots, and the one directly relevant to their interfaceability. Quite recently for instance, 57% end-to-end efficiency at a 1 GHz rate has been achieved with a conventional, gated quantum dot strongly coupled to an optical microcavity [263]. Attempts at similar schemes in Rb-like quantum dots are not as far along in terms of the figures of merit, but microcavities have been successfully used in such dots as well, increasing the extraction efficiency 10-fold in [264], for instance.

I think there is good reason for optimism with regards to the efficiencies of Rb-like quantum dots reaching par with those of conventional dots. Many techniques from conventional dots have seen direct transferal or adaptation to Rb-like ones. Take electrical control, which has long been implemented in conventional dots [265]: both for tuning the emission frequency and to help achieve transform limited spectral performance which crucially depends on overcoming noise [266]. Until recently Rb-like dots usually relied on strain to achieve frequency tuning [267], but transform limited performance was sometimes found under such conditions [268, 269]. This was not the typical result across varying samples however, and dots often had problems with blinking. Recently, however, the gap has closed with a Rb-like dot embedded in an n-i-p-diode [270], granting both frequency tuning by gate voltage and reliable spectral performance by effectively eliminating charge noise. Similarly, some control over the emitted waveform by shaping the excitation have been reported in conventional dots for longer [271], but an implementation in Rb-like dots using Raman scattering also shows convincing shape control [272]. In a nutshell, there seems to be no reason for techniques based on photonic or semiconductor engineering, or quantum optical schemes, not to work just as well in Rb-like dots.

This leaves open the questions of indistinguishability and state-accuracy, and here Rb-like quantum dots shine particularly brightly. The best, current state accuracy result for single-photon sources generally,  $g^{(2)}(0) = 7.5 \times 10^{-5}$ , was the product of a Rb-like dot [273]. In conventional dots, the photon indistinguishability from one dot can exceed 99% [274] and can be improved further by Purcell enhancement though placing the dot in a cavity. Implementing this technique, a single, multiplexed, conventional quantum dot achieving 93% indistinguishability has been used to run  $N = 5$  boson sampling [275] (more on this application in section 2.3.3). Measured indistinguishabilities from remote conventional dots are usually low [276]. Cavities have also been used to improve this figure, but generally still yielded values below 50% [277]. A phonon assisted

technique in remote Rb-like dots managed to just beat this threshold a few years ago [278]. Since then, near unity indistinguishability from a single Rb-like quantum dot was found shortly after the above state accuracy result [279], and has now uniquely been found to be achievable with photons from remote dots (93%), even without cavities, due to the extremely low noise levels obtainable by gating [280].

## 2.2.4 Spontaneous Parametric Down-Conversion

Spontaneous parametric down-conversion (SPDC) is a  $\chi^{(2)}$  (three-wave mixing) process by which a high frequency photon from a pumping laser is converted into a pair of lower frequency photons in an optically non-linear crystal [230, 231]. It is possible to make this process occur degenerately, but this requires greater control and offers no advantages in our envisioned scenario so I will not discuss that case further. We consider therefore the conversion of a pump photon into two distinct photons, labeled signal and idler, related in frequency so as to conserve energy,

$$\omega_{\text{pump}} = \omega_{\text{signal}} + \omega_{\text{idler}}. \quad (2.13)$$

The wavelengths accessible by this technique extend from the visible to the mid-infrared, and sources can be tunable over 100s of nm [281]. The simultaneous conservation of momentum, essential to making the process occur with meaningful probabilities, may need to be enforced with a trick. This is because frequency dependent dispersion ( $n = n(\omega)$ ) in the non-linear medium can lead to a wavevector mismatch  $\Delta k$ . One solution among many is periodic poling, that is, alternating the crystal orientations with a period  $\Lambda = 2\pi/\Delta k$ , producing so called “quasi phase-matching”. The typical length scale of  $\Lambda$  is on the order of 10  $\mu\text{m}$ . As the refractive index also depends on polarization (crystal axis direction), how the phase-matching condition is fulfilled determines whether the signal and idler photons have parallel (type-I) or orthogonal (type-II) polarizations. Crystals prepared in this manner are commercially available, as are entire turn-key sources for a handful of applications. What pumping frequencies will produce the desired signal and idler frequencies can be modeled using Sellmeier equations [282]. These conservation laws are the origin of the strong correlations in the produced photon pairs across their degrees of freedom.

SPDC sources are comparatively simple photon sources in terms of their experimental complexity. When pumped continuously they are probabilistic sources, requiring the detection of an idler photon to herald the presence of a signal photon. When rates are of lesser concern, it is possible to pump the crystals with pulses producing pseudo-deterministic output.<sup>10</sup> Moreover, their properties are amiable to tailoring, and

---

<sup>10</sup>In an interfacing experiment it is then possible to treat the source as deterministic using the pumping pulse as a clock and post-selecting the experiments in which an idler photon was detected.

### 2.3. Overview of Applications

---

although there are always trade-offs, the relations between different figures of merit are predictable enough that a meaningful design optimization can be performed. As a baseline, let us first consider bulk crystal experiments, where the down-conversion process can produce very high heralding efficiencies,  $\eta_h > 95\%$ , and sufficient brightness that multiplexed sources achieve multi-photon rates similar to those achieved with multiplexed quantum dots – compare [283] to [275]. Unfortunately photons from such sources have bandwidths of a few nm and are thus not suited for interfacing with atoms. To narrow the emission without discarding most of the photons in an absorptive spectral filter, the non-linear crystal can be pumped within a cavity resonant to the desired signal and idler frequencies [284, 285]. This way sources at various useful bandwidths have been achieved, be they narrower than 1 MHz [286], a few megahertz [287], or 100 MHz [288]. External cavities around the non-linear crystal lead to losses at the surfaces and wavefront deformation that limit the coupling efficiency to single mode fibers. This can be addressed with a monolithic cavity approach, i.e. polishing and dielectrically coating the crystal itself. As performance then depends sensitively on the quality of the polish and coatings it is difficult to make generally valid statements, but I will discuss a specific case in chapter 4.

The state ideally produced by SPDC sources is a two-mode squeezed state of the form

$$\rho_{ab} = (1 - p) \exp\left(\sqrt{p}\hat{a}^\dagger\hat{b}^\dagger\right) |0, 0\rangle \langle 0, 0| \exp\left(\sqrt{p}\hat{a}\hat{b}\right), \quad (2.14)$$

where  $\hat{a}$  and  $\hat{b}$  are the bosonic operators of the two modes and  $p$  is the pair generation probability. Expanding the exponential in the limit of small  $p$  yields

$$\rho \approx (1 - p)^2 |0, 0\rangle \langle 0, 0| + (1 - p)p(|1, 1\rangle \langle 0, 0| + |0, 0\rangle \langle 1, 1|) + \mathcal{O}(p^2), \quad (2.15)$$

whereupon the detection of a photon in one of the modes (heralding) traces that mode out of the state. This equation immediately reveals the inherent trade-off between state accuracy and rates. The pair generation probability is determined by the pumping power. This sets both the rate of pairs and the rates of multiple pairs, which add a multi-photon component to the heralded single photon. It is also possible for the conservation laws of the process to be compatible with multiple signal/idler frequencies simultaneously, which results in multi-mode emission. In this case spectral filtering of the idler photon can be used to produce a heralded single mode.

## 2.3 Overview of Applications

Ultimately, the role of both single photon sources and quantum memories for photons in the broader technological realm is that of device primitives. Together they can generate,



buffer, retime, and facilitate the coherent manipulation of photonic signals. When memories and photon sources are called enabling technologies, those are the things enabled. As this equates to basic control over the signals, the devices are important generally. A review with an eye on applications summarized quantum memories and photon sources, along with frequency converters, quantum random number generators (qRNG), and photon detectors as the essential components underlying the construction of photonic quantum technologies [50]. Nevertheless, there are plenty of examples where interfaced sources and memories already constitute the totality of the useful device, so it is possible to get a little more concrete. In particular, I will touch on some use cases from the fields of communication and information processing.

### 2.3.1 Quantum Communication: Infrastructure Projects

Quantum communication encompasses everything involving the transfer of quantum states between places [289]. The underlying methods are prepare and measure protocols [290] using pairs of conjugate states, and entanglement protocols based on quantum state teleportation [291]. The former method assumes that the created states are known and is thus dependent on the apparatus, while the latter approach is device independent. Generally, the resource that enables this kind of quantum communication is entanglement shared at a distance. In transmission this typically means entangled photons, although the entanglement may be shared between other objects at other points in the protocols. Connections between two systems have become quite flexible, with a proof-of-principle experiment demonstrating just how heterogeneous connections could be and still transfer states with high conditional fidelity a few years ago [256]. Therein a cold Rb ensemble is interfaced with a rare-earth ion doped crystal via telecom fiber, frequency converting twice. Although the overall efficiency was quite small due to losses, the only significant fundamental loss identified originates in silica fiber absorption at the theoretical limit. At the state-of-the-art, interfacing tools have started functioning, and they will be optimized into greater functionality over time, but even then the underlying problem will not be solved. There is always loss, and due to these channel losses, as well as the impossibility of classically understood signal amplification (no-cloning), the effectiveness of entanglement distribution necessarily scales poorly with the remoteness of the involved parties. The best possible performance for point-to-point quantum links has been quantified in terms of fundamental limits [292]. Both terrestrial, fiber-based links [293] and satellite-based links in space [294] ultimately find their rates limited by loss over distance. To improve upon these limits, quantum repeaters [295] must be used.

The field of quantum repeaters has many facets, including all-photonic protocols, i.e. repeaters without memories [296]. That said, simple repeater protocols pursuing the

goal of improving over direct transmission with experimentally accessible techniques overwhelmingly rely on linear optics and atomic ensemble based memories working in concert [52]. Minimally at the nodes of the communicating parties, memories are conceptually high indispensable to enable further processing. The DLCZ protocol [254] is probably the most influential proposal in this domain. The basic idea of this kind of repeater implementation can be summarized as follows: Consider a network consisting of  $N$  nodes, where the distances between the neighboring nodes are small enough for effective point-to-point entanglement distribution, but too great for viability thereof between non-neighboring nodes. Entanglement is first created between the short distance links for which the process is still efficient. The entanglement is stored in local memories until all the nodes in between the two remote nodes wishing to communicate have established shared entanglement with their neighbors. Local entanglement swapping [297] operations at the intermediate nodes then distribute the entanglement over the full distance. A particularly efficient repeater architecture relies on single photon sources and compatible quantum memories [51]. Therein each node consists of a deterministic single photon source and a quantum memory. Entanglement is established between the nodes  $A$  and  $B$  by having each party generate a photon and send it through a beam splitter. At both nodes, one mode after the beam splitter is stored in the local quantum memory, while the other is sent to a remote central station for a Bell-state measurement. Upon detection of one photon at the central station, the resulting state consists exclusively of a vacuum term and a term proportional to the entangled memory state  $|\Psi\rangle = \sqrt{2}^{-1}(|1_A 0_B\rangle + |0_A 1_B\rangle)$  representing one photon stored in one memory. This approach avoids an error term corresponding to two full memories present in schemes using pair sources, which competes against operation at a high rate. A real system has to optimize the ratio of its beam splitter for losses to maximize the success rate. Later, to perform Bell tests with path entangled states, a displacement operation can be used as measurement setting [298, 299]. A parallel approach to high repeater rates pursues multi-mode operation [148]. The effective difference is that multi-mode repeater protocols can be attempted with every mode simultaneously, dividing the time it takes to establish entanglement by the number of modes. Multi-mode approaches are highly varied and include spatial multiplexing [300], wavevector multiplexing [301, 302], and spectral multiplexing [303].

There is a simple way to summarize what the quantum repeater does – it uses a memory to overcome a timing problem. The core on which the rest of quantum communication builds is a series of protocols that are practically facilitated by memories in this same way [55]. Entanglement purification [304] is another example of a scheme to facilitate the establishment of shared communication resources. The goal therein is to produce stronger entanglement by reducing the number of states. For a concrete example let us consider conversion of a large number of weakly polarization-entangled

photonic pair states shared by parties A and B, which might be what they actually end up with after an attempted entanglement distribution using photon pairs, into a smaller amount of approximate polarization Bell states. This *distillation* can be accomplished using linear optics. Each party overlaps their halves of their shared pairs on a diagonal polarizing beam splitter and detects one of its modes [305]. One probabilistically obtainable result, after the parties share their measurement results with each other, is a more strongly entangled pair. This is an example of a protocol using what I previously labeled LOCC, local operations and classical communication. For something like this to work at scale and at a distance, timing control over the photon pairs must be established – after all they need to be precisely overlapped pair-wise. Practically, they will also need to be held while the parties share their measurement results, which is also true for any communication protocol requiring multiple rounds. Most importantly, however, if the parties can store their successes they do not need every round of the protocol to succeed simultaneously, and the improved scaling is immediately evident in implementations [306]. Similar success rate enhancements by using memories apply to diverse LOCC protocols beyond ones facilitating communication, some more examples given in [55] are continuous variable cluster state quantum computation and quantum illumination.

### 2.3.2 Quantum Communication: Cryptography

The field of communication’s theoretical concerns are generally about more than just how to transfer information. Communicating parties may want to enjoy certain guarantees with regards to the information exchange, perhaps that it is secure against eavesdroppers. To be able to make such guarantees cryptographic protocols are used. Encryption to protect secrets is baked deep into cryptography’s history, and so are other protocols designed to produce desirable attributes like anonymity or the certification of identities, records, or even randomness. Activities enabled by these protocols include war, a free press, an electronic economy, airplanes that do not stall and kill their passengers when one of their instruments faults, and gambling. Like most technology, it has its ups and downs. Quantum communication addresses all the same questions and problems that are tackled with classical cryptography, just with the exchange of quantum information. The most significant promise this approach makes is to lay the foundations of security in the laws of physics.

The first schemes of quantum cryptography came in the form of protocols to generate symmetric secret keys [307]. They form a group of experiments categorized by their assumptions on implementation devices called (measurement device independent) quantum key distribution or (MDI)-QKD, which have become increasingly practical in the last decade [308, 309]. Symmetric secret keys are cryptographic primitives that

### 2.3. Overview of Applications

---

can be used in a variety of protocols like encryption and authentication. Secrecy is a relative quality and any quantitative measure of it must consider both the probability of the secret being intercepted and the probability of later decoding. The highest standard of secret is that based on physical security, as the probability of interception can be limited with force and the probability of decoding can be made zero. A way of doing this is for communicating parties to share two code-books, each containing a copy of a one-time pad of random numbers,<sup>11</sup> then maintain a physically secured chain of custody of the code until it is destroyed after use. The involved logistics are difficult, especially in scenarios like communicating with anonymous strangers, and because the code is consumed by the protocol (security is only maintained if it is used just once) it is a resource that requires periodic renewal. Yet such a symmetric code is a resource immediately available through what I just labeled the “quantum communication infrastructure projects”. Using the entangled photon pair established by the means of the previous section, the two communicating parties can generate a key by violating a Bell inequality with their shared state. Their degree of security, or the chance of having been eavesdropped on, is then physically bounded by the violation they observe due to the monogamy of entanglement. The parties can then sacrifice key rate for security via privacy amplification [312] to reach an arbitrarily low chance of having been spied on. There are some variations in the details of different protocols, so this is just an example. The general result is that there are means to securely expand shared one-time pads between parties remotely, after a one time key expenditure to authenticate their classical communication.

The QKD protocol is not invulnerable. A motivated hacker might find a weakness in a real device where the theoretical security breaks down and leaves room for undetectable wire-tapping of the secret key [313]. A brazen hacker might even manipulate the secret key with diverse methods, including damaging the photon detectors of a communicating party with a laser [314]. The ultimate security is therefore necessarily subject to evaluation against the feasibility of attacks requiring a little field work [315]. In security schemes generally, from bike locks to armoring, the outlandishness of the attacks that a system can resist is called its *hardness*. A destructive approach to vulnerability testing is a fruitful way of developing countermeasures, but these kinds of problems are not special to the quantum case. What is different to the case of classical computer cryptography

---

<sup>11</sup>Properly, as its reliable generation is not trivial, randomness itself should be treated as a resource with a cost. To avoid too much of a tangent on the matter, I invite the reader to leave aside concerns about the principally deterministic nature of tossing coins, rolling dice, and spinning roulette wheels, and to imagine the randomness I postulate in this chapter to emerge from their favorite physical process producing outcomes that are unlikely to bankrupt a casino when treated as random. More on the usefulness of random numbers and novel techniques to obtain them can be found in a review on quantum random number generation [310]. Further, a short but interesting section on dealing with imperfect randomness in QKD protocols is contained in [311].

over the internet is that once it is shared, the quantum secret is secured by information theory.<sup>12</sup> All the risk occurs at once and, combining privacy amplification and a risk analysis to the dangers of a physical attack, a tolerance bound can be set. In the cryptographic schemes in common use for securing telecommunication data today, the quality of the secret is evidenced by the computational complexity of a cryptographic algorithm. The typical example demonstrating the dangers inherent to such schemes is RSA public key cryptography, which relies on the computational difficulty of factoring large primes. While it is not feasible to do so yet, this cryptosystem is efficiently broken by a quantum computer running Shor's algorithm [316]. A recent run-cost analysis thinking of superconducting qubit implementations calculates requiring just  $1.4 \times 10^4$  logical qubits to factor 2048-bit RSA in 8 hours [317], estimating  $2 \times 10^7$  physical qubits would be required using some qubit quality assumptions and surface codes for error correction [318]. Naturally, the required assumptions about future qubits put large uncertainties on the exact numbers, but given a couple more decades this scale is at least conceivable. For data with long term value, medical or genetic data, or data legally requiring protection like census data, this is a relevant security risk today as the encrypted data could be harvested now and decrypted in the future. This is not to imply that all standard cryptography based on computational complexity is doomed to fall to quantum computers [319], even if they are scaled up along the lines of the most optimistic predictions. For many kinds of problems quantum computers have no known (advantageous) algorithms or only quadratic advantages over classical algorithms,<sup>13</sup> which puts into doubt that they are a more practical approach to cryptanalysis generally [320]. Computational schemes with resistance to quantum cryptanalysis [321] in mind are called post-quantum cryptography [322] and the National Institute of Standards and Technology is well underway to standardizing quantum-resistant public-key cryptographic algorithms [323]. These observations are not at odds with each other because of the difference in the underlying security mechanism and required implementation effort. Once again we arrive at a case of different tools for different jobs, practically to be selected by risk or cost-benefit analysis.

There are many further quantum communication protocols, usually carrying the promise of information-theoretic security. Examples include authentication, which is a guarantee that a sender is who they say they are. If the guarantee is transferable, then

---

<sup>12</sup>Information-theoretical security is also granted by a cipher made from a physically exchanged one time pad and is sometimes called “unconditional” security. This term was born out of contrast to the security provided by the expected maximum expenditures an attacker could bring to bear on breaking a key, i.e. it refers to security proofs willing to assume an attacker has infinite computational resources. Given a holistic view, it is surely a poorly chosen term for any real system, notwithstanding underground bunkers. In [308] an explicit list of the conditions on “unconditional” security is given, and the impossible concept of a security beyond all conditions is labeled “absolute” security.

<sup>13</sup>At <https://quantumalgorithmzoo.org/>, a comprehensive catalog of quantum algorithms is maintained.

it is a signature. Secure quantum authentication of a message requires communicating only one qubit [324], and quantum signatures implemented directly [325] can compete with signatures generated from QKD produced shared key. This is always a good baseline to compare a specific quantum protocol against, as any cryptographic protocols based on a symmetric key obtained from QKD in this way inherit its security. In a similar vein of tasks, fingerprinting is a protocol used by two parties to anonymously verify that they belong to the same group. The quantum protocol is exponentially more efficient in information communicated than is possible classically [326] – a bound violation that has been demonstrated [327]. These efficient results do not necessarily imply that these protocols will replace their classical counterparts, just that they are cheap in a quantum network.

Practically scaling protocols to larger numbers of involved parties is also investigated. Recently an anonymous and secure four-party quantum communication protocol could demonstrate key rates around  $10^{-2}$  bit s<sup>-1</sup> [328]. An advantage of anonymous quantum protocols is once again information-theoretic security [329]. This avoids the common classical work-around of parties trusting their privacy to commercial proxy services or virtual private networks, or even to distributed proxy networks like TOR – more on decentralized systems in the final section. Whether communicating parties can trust one another is also a recurring delineation in protocols, particularly as to their difficulty. Bit commitment is a prime example. This is a protocol that serves as a primitive for coin-tossing between distrustful parties and two-party secure computation. A no-go theorem proves that it is impossible to protect the envisioned protocol against cheating information-theoretically [330, 331], even with quantum cryptography. However, a later extension showed that quantum methods could bound the cheating [332], and of course protocols under various security conditions were investigated as well. More theoretical protocols seem to circumvent the no-go theorem by using quantum states not covered under its assumptions [333]. For a more detailed overview of the landscape of implementations, the recent review by Xu et al. [309] contains expansive tables of experimental demonstrations of QKD and other protocols, including hacking. Typical key rates of implementations listed therein, which depend on the quality of sources and memories used even if the security does not, currently rarely exceed a few kbit s<sup>-1</sup>. While this sounds a lot like quantum dial up, it represents an increase of at least four orders of magnitude in a decade, and the road to Mbit s<sup>-1</sup> key rates is steadily being paved by increasing clock rates.

#### 2.3.3 Quantum Computing Optically

Optical quantum information processing (QIP) has been known to be possible in principle for a long time. Optical models of the universal and reversible Fredkin, or

controlled swap, gate have been known since the eighties [46], but require a cross-Kerr nonlinearity. Due to a confluence of weak interactions and noise, implementations in bulk material still seem just out of reach over 30 years later [334]. Instead proposals for nonlinear gates using atoms in cavities [335, 336] or atomic ensembles [337, 338] have gained some traction, and often resemble previously discussed memory systems [338, 339]. In the year 2000 two influential methods for linear optical quantum computing were proposed. The circuit scheme uses beam splitters and phase shifters for operations, also requiring single photon sources and detectors [44]. The measurement-based scheme [45] uses entangled states called cluster states [340] and runs programs via detection. These schemes are conceptually different yet have numerous similarities. Both gain the nonlinearity required for universal computation through photon measurements, and both are probabilistic, with similar resource costs scaling their success rates [341]. Indeed efficient scaling of probabilistic gates generally relies on repeater like protocols and heralding [342, 343]. In all scenarios, photon-atom interfaces can play important roles [218, 344].

Due to efficiency requirements, many single photon sources used today are heralded. These devices are at the heart of optical QIP, yet they, and the computation itself in these linear schemes, are both probabilistic. A synchronization solution to the implicit scaling problem for heralded photon sources using memories has been analyzed in detail by Nunn et al. [59]. The rate at which  $N$  heralded photon sources, each with a photon heralding probability  $p \ll 1$ , can generate a state of  $N$  coincident photons is  $c_{N\text{-coin}} = p^N$ . In [59] the authors suggest equipping each source with a memory, capable of storing its output photons with an efficiency  $\eta$  and for a lifetime  $\tau$ . Assuming that the output rate of the sources is limited by the photon bandwidth, and this in turn is matched to the acceptance bandwidth  $\delta\nu$  of the memories, then the  $N$ -fold coincidence rate is improved to  $c_{N\text{-coin}} \approx (p\eta\tau\delta\nu)^N$  by the synchronization procedure. This scenario highlights the importance of the time-bandwidth product  $B = \tau\delta\nu$ . Synchronization has begun entering the business of enhancing multi-photon rates [78, 345], but not at scale and not with systems having base rates comparable to the multiplexed systems discussed previously in section 2.2.4. A state of  $N$  coincident photons has an interesting direct application in boson sampling [346]. A boson sampler is an interferometric circuit of beam splitters and phase shifters that accepts such a state as input and measures (samples) the output. While computational problems can be encoded in the circuit, this analog computer is not believed to implement universal QIP or even classical universal computing. Simultaneously, however, classical simulation of the circuit would require calculating the permanent of a  $N \times N$  matrix with complex-valued coefficients, which is unreasonably hard.<sup>14</sup> Implementations have found success with laser written

---

<sup>14</sup>The authors doing the analysis in [346] claim, that it is easy to show that even approximating this permanent is a  $\#\text{P}$ -complete problem. The evident motivation of constructing such a system is testing

circuits [347], and the results of classical simulations can be used as benchmarks [348]. Using squeezed light as input, an approach called Gaussian Boson sampling [349], an experiment with 50 identical squeezed inputs could recently be performed [350]. The rates observed show a quantum advantage present in this result, but the comparison depends on the classical simulation it is compared to. Where the ultimate boundary therein lies is still being investigated, but the demonstrated advantage does not seem to be as overwhelmingly large as initially claimed [351]. A recent, more detailed overview of boson sampling and other optical QIP implementations can be found in [218].

#### 2.3.4 Quantum Computing on a Quantum Network

No physical system attempting to implement quantum computation currently investigated is free from technological limitations to its scaling. In many cases it is decidedly non-obvious that it will simply be possible to *cram* more qubits on to chips akin to what Moore divined for classical transistors [3]. Open “engineering problems” are diverse across systems, from the space on a microchip to the thermal load induced by control and read-out a dilution refrigerator can bare. Moreover, comparisons between implementations are difficult because they ultimately must take computer architecture into account [352]. The commonly envisioned solution to whatever scaling limitations a single device type will ultimately encounter comes in the form of distributed computing [353], and an interfacing standard has been proposed [354]. The idea, also called telecomputation [355], is to perform calculations on a network of linked quantum computers of whatever size is viable. A high level view of the layers such a scheme requires is given in [356]. An important note is that in a scenario where processors are being connected for scaling the computation, there is no inherent need for the network connecting them to be particularly remote. The speed and interfaceability of the communication devices would be paramount, but operational distances could be kept to local area network scales. This would mitigate the issue of loss and the importance of working at the telecom wavelength, as well as loosening the demands on lifetime (but not on time-bandwidth product [60]) which would only need to be long in comparison to the timing of logical gates, and not necessarily long in comparison to a communication latency over distance. This opens up the field of systems directly relevant to implementing networking without frequency conversion considerably.

A particularly interesting problem in distributed systems is that of getting the components to agree on a shared reality relevant to the computation, i.e. having them form a consensus. The issue arises in the presence of active sabotage, but also when faults in the subsystems or communication channels lead to conflicting information

---

the extended Church-Turing thesis without building a full quantum computer. The thesis states that a probabilistic Turing machine can efficiently simulate any realistic model of computation. A real computing system intractable to classical simulation would prove the contrary.



being sent around by a single node. An abstract description known as the Byzantine generals problem is impossible to solve unless at least two-thirds of the networked system functions properly, and insurance against a large number of faults, called Byzantine fault tolerance, is classically expensive in required communication rounds [357]. These are important considerations when lives are entrusted to computers, like with autopilots, or for networks deriving a benefit from the consensus, like Blockchains or TOR. Quantum Byzantine agreement protocols using high dimensional entanglement can detect opportunities to form a consensus before attempting a broadcast [358], can then immediately reach the more intuitive limit of consensus in the presence of any saboteur or fault minority [359], and can operate in constant time [360]. These are difficult problems with significant quantum improvements over classical resource costs [361]. An established consensus can detect faults or saboteurs and repair or exclude their nodes as applicable [362]. Small scale quantum schemes have been implemented using four photon entanglement [363]. There are a few further approaches for defining consensus in a quantum network, but summarizing literature is scarce. A brief overview may be found in [364].

In closing this section I want to summarize the most relevant application context for the ultimate experimental results. A heterogeneous photon source-memory interface is presented in chapter 5. This system is particularly interesting with regards to the applications requiring fast operation – local networks, synchronization, and QIP. The memory implementation in hot vapor represents a relative simplification to other systems, which opens up scalability via the proven route trod by vapor cell based atomic clocks once the design is matured. The final time-bandwidth product demonstrated is  $B = 250(20)$ , with no significant efforts made to improve the lifetime beyond the simplest limits identified in section 2.1.4. Indeed, were a photon synchronization experiment as described in [59] attempted with it, some multi-photon rate enhancement would be expected, although it would be tempered by the efficiency. This is despite the fact that the experiment is a proof-of-principle of photon storage in hot vapor ground-states and the first demonstration of a highly non-classical  $g^{(2)}$  upon readout in comparable systems. Simultaneously, it is important to acknowledge that the requirements imposed on photons with regards to state accuracy, indistinguishability, and efficiency for advanced applications like repeater networks or QIP are high [229], and that they are not currently simultaneously met by any single system or interface.



## Chapter 3

# Hyperfine Memory

*The first principle is that you must not fool yourself  
and you are the easiest person to fool.*

Richard Feynman

---

This chapter discusses benchmarking experiments of a lambda memory implemented with the hyperfine states of the  $^{87}\text{Rb}$   $D_1$  line, characterized with weak coherent pulses. The ultimate motivation behind this work is the goal of interfacing single photons from solid state sources with an atomic memory in hot vapor. The main results are published in [365],<sup>1</sup> and are elaborated upon here. Many of the necessary improvements and modifications to make a memory truly suitable for storing single photons were elucidated by these results, and they served as a guideline for the necessary performance of the photon source. In particular, we envisioned an interface with quantum dots facilitated by an in-house collaboration. The state of these sources at the time is well captured by the contemporaneous publication [272]. After these experiments in the hyperfine memory, we concluded that more work was necessary to make the systems compatible. This motivated the construction of the SPDC photon pair source tailored to our hot Rb memory introduced in chapter 4, as well as to the revisions of the memory that enabled their interfacing described in chapter 5, but let us start at the beginning.

---

<sup>1</sup>A note on credit: These experiments were performed in the first year of my doctoral studies. Janik Wolters began work on the project before me, including some preliminary measurements which I do not describe, and in particular he did the four-level theoretical modeling. The experiments I do describe were set up and performed together with him. After five months we had characterization results we thought were worth publishing, so I analyzed the data and he began writing the letter, reconvening and collaborating on revisions and the review process. As such, the first-person plural in this chapter generally refers to Janik and me, unless a wider context is obvious.

### 3.1 Scheme and Implementation

The  $^{87}\text{Rb}$   $D_1$  line between the electronic configurations  $5^2S_{1/2} \leftrightarrow 5^2P_{1/2}$  consists of two hyperfine split ground states and two hyperfine split excited states. Herein the total spin takes on the values  $F \in [1, 2]$  due to this isotope's nuclear spin  $I = 3/2$ . Including the Zeeman degeneracies, this means the line consists of a total of 16 atomic states, with polarization dependent transition strengths between them where dipole transitions are allowed. The dipole matrix elements of these transitions are tabulated in [115], and the relative strengths for all alkalis are diagrammed neatly in [366]. Nevertheless, even armed with this information, this generates a lot more structure than we encountered in the theoretical models of chapter 2. In the case of  $\pi$ -polarized light, for example, the selection rules allow transitions between states with  $\Delta F \in [0, 1]$  and  $\Delta m_F = 0$ , unless  $\Delta F$  and  $m_F$  are both zero. Lambda systems formed with signal and control in orthogonal polarizations are of particular interest for the purpose of combining and separating these modes before and after the memory, respectively. An exercise in counting for crossed linear polarizations reveals 15 possible two-level transitions which form 6 lambda systems,<sup>2</sup> see figure 7.1 in the appendix for an explicit diagram.


Recalling that EIT itself is a phenomenon of interfering absorption pathways that occurs with just three levels, it lies close at hand that further interference occurs and modifies the situation here. Due to the polarization dependence of the matrix elements, in particular with regards to their signs, this is also the kind of analysis required to formally determine which signal and control polarizations will even form a lambda system. Having acknowledged this, let us consider the effective implications for experimental implementation. Wanting to use the hyperfine states as they are, let us assume that the Zeeman levels are neither resolved nor prepared, i.e. after state preparation the atoms all occupy one of the hyperfine states but remain equally distributed among that state's Zeeman levels. Seemingly, we can immediately reduce our analysis to the level of the effective dipole moments of a four-level system. The only remaining, untreated source of potential destructive interference lies between the two lambda systems involving different excited states. I will elaborate on the boundaries of the case of significant destructive interference in section 5.1.1. For now, suffice it to say that if the hyperfine states are well resolved a good choice of the detuning from the excited states,  $\Delta$ , should ensure that one pathway dominates. Moreover, interference in the absorption pathways does not destroy the signal. Its effect can be understood as a reduction in interaction strength and compensated for as such, e.g. by increasing the optical depth through heating the atoms. For any chosen  $\Delta$  the effective strength can be calculated.

---

<sup>2</sup>Fifteen and six, rather than eighteen and eight, due to those forbidden transitions of the form  $\Delta F = 0$ ,  $m_F = 0$ , ominously foreshadowed by this footnote.

The four-level system we are left with following this simplification is illustrated in figure 3.1. The choice of polarizations, vertical for the signal and horizontal for the control, is motivated by the better polarization extinction ratios of polarizing cubes for horizontal light, polarization filtering of the strong control in mind. The chosen roles of the ground states is the only sensible one, as if the signal were coupled on  $F = 2 \rightarrow F'$  many excited states that could absorb the signal would not couple to the storage state. In contrast, this system appears to enable full EIT. These considerations, including what I have said about interfering systems, were analyzed when the idea of EIT memories was still very young [367], nevertheless setting the blueprint for these choices. The catch, and contrast to [367], is that the atoms here are hot. While the excited state hyperfine splitting of 814.5 MHz may seem fairly large, it is in fact on the same order as the Doppler broadening in Rb near room temperature, so it is not possible to exclude any of the four levels from consideration even for narrowband signals. We are forced to conclude that even in this scheme there is an excited state at least off-resonantly available to the signal which is not addressed by the control, namely  $|F' = 2, m'_F = 0\rangle$ . This is due to the aforementioned forbidden nature of transitions with  $\Delta F = 0, m_F = 0$ . In other words, the signal can be absorbed without being stored. The problem is perhaps more obvious when considering the Rb D<sub>2</sub> line where the excited state total spin can also be 0 or 3. Neither the  $F' = 0$  nor the  $F' = 3$  excited states form lambda systems between the ground states as this would require  $\Delta F > 1$ , but both contribute to absorption. Unlike the analogously forbidden coupling of atoms beginning in  $|F = 1, m_F = 0\rangle$  to  $|F' = 1, m'_F = 0\rangle$ , or even interference between different paths of the storage process, the existence of a purely absorbing transition cannot be reasonably compensated for in general terms. Indeed, it fouls the principle idea of simply “cranking up” the memory efficiency by employing ever higher optical depth, as it represents a photon loss channel that, by definition, scales exponentially with the OD. For this reason we call absorbing transitions of this kind *parasitic single-photon transitions*. When they are present, the working point of the memory  $\Delta$  and atomic temperature must be chosen so that their optical depth remains small.

Let us now turn our attention to the required components to practically implement a memory based on this scheme, targeted at the storage of broadband light from solid-state single-photon sources. As discussed in section 2.2, one key advantage of such sources are the rates and associated bandwidths. These however, are also less naturally compatible with the typical paradigm of atomic physics, especially cold atomic physics where the long trap and coherence times mean experiments neither require nor particularly benefit from fast devices. To bridge the gap between the convenient and slow CW diode laser and the fast time scales at play one possible option is electro-optic modulation. Originating in a context where a microsecond is reasonably fast, almost every single timing related element relevant to the memory experiments



**Figure 3.1:** Hyperfine resolved, i.e. four-level, energy diagram of the  $^{87}\text{Rb}$  D<sub>1</sub> line and transitions involved in the memory experiments, omitting the 3-fold (5-fold) Zeeman degeneracies of the  $F = 1$  ( $F = 2$ ) states. In the notation of chapter 2 we have  $|g\rangle = |F = 1\rangle$  and  $|s\rangle = |F = 2\rangle$ . The vertically polarized signal to be stored is detuned by  $\Delta$  from the  $F = 1 \rightarrow F' = 1$  transition, while the horizontally polarized control laser is detuned by the same  $\Delta$  from the  $F = 2 \rightarrow F' = 1$  transition, ensuring two-photon resonance. The values of the hyperfine splittings are taken from Steck’s collection of alkali D line data [115].

discussed in this thesis was eventually replaced. One choice we stuck with is to generate optical pulses with integrated waveguides of lithium niobate ( $\text{LiNbO}_3$ ) crystals. There are downsides to this approach – some of them become apparent in our experimental results, and sometimes when they limit what is achievable they are reported on in the broader literature [203]. Simultaneously they impart the full flexibility and much of the capability of high-bandwidth electronics directly onto light at a fraction of the cost of less naturally flexible, pulsed light sources. In particular, they make it possible to generate the fast control pulses required to store photons with 100 MHz–1 GHz bandwidth, as well as laser pulses that resemble the waveforms of the targeted photons in their shape and bandwidth. Why do we even need the latter though?

Even if a single-photon source is readily available, single photons are simply too small an amount of signal for the practical construction and maintenance of an experimental setup. Therefore a simple, controllable signal test light source, a laser, is definitely required, be it just for the purposes of spatial alignment on the free-space scale of an optical table. Due to the wide and simple frequency and intensity tunability, as well as standard laboratory practices to greatly ease its operation such as Doppler-free atomic reference spectroscopy, it would be nice if this laser could serve as a passable substitute for actual single photons from solid state-sources for the purpose of benchmarking the performance of the memory directly. This requires two things. First is pulse shaping to match the photon bandwidth as just discussed. Second is a method of calibrating

the pulse intensity. Since lasers are sources of coherent light, the best we can hope for is pulses that contain one photon on average, conveniently thought of as a coherent state with amplitude  $|\alpha|^2 = 1$ . This calibration is essential for the validity of efficiency and signal-to-noise measurements when testing the memory.

Generating the control pulse requires nearly the same modulation capabilities as the signal, but with the opposite power requirements. Recalling equation 2.3, the acceptance bandwidth of the memory depends on the control pulse Rabi frequency, practically its intensity. To achieve sufficiently high intensity with a reasonably sized beam diameter, amplification may be required. Here the qualifier “reasonably sized” is technically a function of the vapor cell. Indeed specialty micro-fabrication of vapor cells has played such an important role, for instance, in the technological development of atomic clocks that it is easy to imagine advances in fabrication techniques revolutionizing most elements of hot vapor physics. For memories, the important role of the optical depth increases the difficulty of this task. The kind of geometry that is truly desirable consists of a long and narrow cell, akin to the  $300\ \mu\text{m} \times 300\ \mu\text{m} \times 10\ \text{mm}$  square channels used both in memory [368] and opto-mechanical [369] experiments of the Polzik group. However those vapor cells are the result of long standing collaborations with academic experts of fabrication, still fairly short compared to typical, commercially available spectroscopy cells, yet still not quite narrow enough that a diode laser would not require amplification to achieve sufficient Rabi frequencies for the storage of photons with bandwidths of 100s of megahertz. We therefore opt for control pulse amplification and commercially available vapor cells.

This covers three essential components, the signal, the control, and the atoms. The remaining requirements are pumping, i.e. preparation of the atoms in  $|F = 1\rangle$ , and filtration of the intense control and induced noise from the single-photon-level signal. I will now describe these components beginning with the pulse generation separately before turning to the optical elements of each identified component. This division of the setup into components is quite natural not only due to the separate purposes, but because for the most part they congruently separate into self-contained and independently optimizable building blocks connected by fiber links.

## 3.2 Pulse Generation

The natural waveforms of single photons from quantum dots rise sharply and fall exponentially. From SPDC sources both the rise and fall of the waveforms are exponential due to the heralding. The shapes of fully optimal control pulses depend on the signal and can be complicated. Shape-optimizing simulations using our targeted quantum dot photon waveforms as input revealed that in our case fully optimal control pulses most closely resemble modulated log-normal distributions. Fortunately, as I will show later,

optimization of Gaussian control pulses yields nearly the same memory efficiency, and indeed optimization under a Gaussian constraint is common practice for broadband memories [370]. Assuming near transform limited light, the targeted bandwidth range of 100 MHz–1 GHz requires pulses with widths on the timescale of 1 ns, and a little more temporal resolution would surely help with accurate shaping. In this experiment, these pulses are generated from continuous wave diode lasers by means of electro-optic amplitude modulators (EOMs) as they can achieve these speeds. The required voltage pulses are generated with a fast arbitrary waveform generator (AWG), which also serves as the experimental master clock.

### 3.2.1 Electro-Optic Modulators

Advancing rapidly to meet growing bandwidth demands, integrated LiNbO<sub>3</sub> modulators had established themselves as the go-to modulation devices for telecom applications by the turn of the millennium [371]. While the demand for high quality devices at non-telecom wavelengths is on a smaller scale, they have also become commercially available in the intervening years. The principle of operation is simple. The refractive index of LiNbO<sub>3</sub> changes when a bias voltage is applied to it by the Pockels, or linear electrooptic, effect. Simply put, the phase of light guided in LiNbO<sub>3</sub> can be controlled electrically. A Mach-Zender interferometer (MZI) made of LiNbO<sub>3</sub> can thus be used to modulate the amplitude of light passing through it by applying a voltage to an interferometer arm. To achieve high modulation frequencies and bandwidths, a fiber-connected, integrated design of these LiNbO<sub>3</sub> MZIs is the standard. Waveguides of the optically non-linear material are fabricated in wafers, whereupon electrodes are directly deposited on the surface and patterned by photolithography. This extreme proximity of the electrodes eliminates the need for high voltages, facilitating operation at high speeds. To further halve the required voltages, typically an equal and opposite phase shift is induced in each of the interferometer arms by arranging the electrodes in a “push-pull” configuration. The voltage required to induce a phase shift of  $\pi$  is called the half-wave voltage  $V_\pi$ , and although manufacturers generally try to make their devices transparent at zero bias the maximum transmission is usually achieved at an offset voltage  $V_0$ . In waveguides typical values of  $V_\pi$  are a few volts, while in bulk they are commonly 3 orders of magnitude larger. The optical power after an amplitude EOM can be stated using its transfer function,

$$P_{\text{out}}(V) = P_{\text{min}} + (P_{\text{max}} - P_{\text{min}}) \left( \frac{1}{2} \cos \left( \frac{\pi(V - V_0)}{V_\pi} \right) + \frac{1}{2} \right). \quad (3.1)$$

The maximum output power,  $P_{\text{max}}$ , is determined by the power available before the EOM, subject to the device’s damage threshold, and the insertion loss which is typically



a few dB. Also note that LiNbO<sub>3</sub> waveguides can be polarizing depending on waveguide type. As ours are, this statement assumes that the light polarization is correctly aligned. The minimum output power,  $P_{\min}$ , in turn is a function of the EOM's contrast or extinction ratio.

In all our experiments we use Jenoptik AM785 modulators. These devices tolerate around 50 mW of optical power, and biases somewhere north of the cautiously specified  $\pm 10$  V. Photorefractive damage, which LiNbO<sub>3</sub> is prone to when guiding light shorter than about 1  $\mu\text{m}$ , accumulates over time, so operation at optical powers at the upper end of the tolerances will eventually degrade the EOM. This manifests itself as greater insertion loss and lower contrast than initially achieved. There is some variation in the performance from device to device. Measured insertion losses fall in the range of 5 dB–6 dB, and extinction ratios are generally around 30 dB or slightly lower. When measuring the extinction ratio it is essential to test the EOM in a time resolved manner while generating pulses; the ratio measured when applying a DC bias manually (often  $> 40$  dB in these devices) can significantly overestimate the achievable performance a short time after a pulse is generated due to ringing. Browsing the catalogs of various manufactures of similar devices it is easy to find claims of significantly better performance. From past experiences with EOMs before I began memory experiments, I had already concluded *caveat emptor*, or let the buyer beware, when it comes to claims of significantly higher extinction ratios or much lower insertion losses.<sup>3</sup> The test scenarios with which claims of extinction ratios are supported can be as extreme as what is achievable for a single pulse after careful, manual optimization of different bias voltages applied to multiple electrodes. Likewise, insertion losses are sometimes misleadingly specified as just the absorption in the waveguide, without accounting for losses coupling light into the waveguide via the fiber-pigtail. This is a deceptive practice because LiNbO<sub>3</sub> is highly transparent to NIR light, while fiber connections to waveguides are often subject to a large mode-mismatch which inherently limits the coupling efficiency. For more on LiNbO<sub>3</sub> waveguides including details on the in-coupling problem and mitigation strategies see [372].

From the transfer function, the method to generate optical pulses can be read-off. The EOM needs to be continuously biased to the point of minimum transmission at  $V_0 \pm V_\pi$ , after which voltage pulses with an amplitude of  $V_\pi$  can be applied to generate similarly shaped optical pulses with the amplitude  $P_{\max}$ . Experimentally, the

---

<sup>3</sup>I mention this in part due to the contrast to my experiences with our present EOM supplier, who provides specifications for realistic experimental situations and has always been responsive and knowledgeable about what kind of performance we could expect in our use case. Inquiring on the telephone a couple of years after this experiment about whether poorer performance we were seeing was likely to be damage, we were met with frank surprise that the EOMs had lasted this long under our operating conditions, as well as learning about the painful compromises RoHS compliance brings with it.

distortion induced by the cosine to the pulses' peaks and tails are small and not worth attempting to compensate for, possibly due to the limits of electronic pulse generation at high speeds with our devices. As  $V_\pi$  is a function of the pulse repetition rate, and to account for the fact that slightly smaller peak voltages keep the EOM transfer function in a more linear regime which could result in better behaved optical peaks, the amplitude of the voltage pulse should always be optimized by measuring the optical pulses actually generated. This is simple enough, however there are other complications. The bias required to keep the EOM at its minimum-transmission working point for pulse generation drifts. Three main physical effects are at play: temperature variation induced drifts via the pyroelectric effect, photorefractive effects due to the guided light, and longer term transient effects due to the applied bias itself (charges accumulate) [373]. Temperature stabilization reduces these drifts, as does operating the EOM at low optical powers, but neither mitigates the drift sufficiently to preclude the need of active compensation to remain at the optimal working point on the timescale of a few seconds.

#### Locking EOMs to the Optimal Operation Point

When operated at low duty cycles, i.e. when the desired optical pulse width is small compared to the pulse repetition rate, EOMs can be held at the minimum transmission point with lock-in techniques. This approach is used throughout the experiments of this thesis, with one exception which I will explain shortly. I have not performed a systematic investigation into how low the duty cycle has to be, so I can only say that this technique still performed sufficiently well at the largest duty cycle investigated, which consisted of a 1 ns wide exponentially falling signal pulse (see section 3.3.1) generated every 400 ns. The technique is implemented as follows: After the EOM a beam sampler, i.e. a glass plate without dielectric coating, picks up 1%–10% of the light by Fresnel reflection and directs it to a photo diode. If the amount of light is insufficient to make use of a regular diode's dynamic range, then it is worth using an amplified detector. In our case the Thorlabs PDA36A2 with up to 70 dB gain has established itself in this role. Its signal is extracted via a lock-in amplifier so that it can be stabilized to its minimum value. We do this digitally with a Labview routine controlling a computer integrated National Instruments card (NI PCIe-6363) connected to a breakout box (NI BNC-2090A) to produce the optimal bias. An inexpensive but nevertheless needlessly capable amplifier (aim TTi WA301) is abused for impedance matching. This slowly varying bias is then combined with the RF pulses on a bias-tee (Mini-Circuits ZFBT-4R2GW-FT+) before reaching the EOM. The drifts are slow enough that lock-in frequencies of a few kilohertz suffice. Lock parameters are optimized with respect to the generated optical pulses. Drifts tend in the direction of the bias, so

after some time the optimal bias approaches the maximal voltage that can safely be applied to the electrodes. For ease of operation our locking routine then automatically jumps by  $2V_\pi$  towards zero bias to target relocking at the next lowest bias optimum of the periodic transfer function. Colloquially we call this a “turnover”. With some tweaking of the parameters, the lock systematically catches this new optimum and the operation is disturbed only briefly. At high optical powers turnovers can occur at rates of one every 1 min–2 min, while at lower powers they are almost never necessary. The relevant timescale of comparison is the experimental integration time. If it is long compared to the turnover rate, then it would be appropriate to note them and exclude the experiments performed before the EOM resettles, but this proved to be unnecessary here.

If operation at low duty cycles is not possible, then the optimal EOM bias must be stabilized some other way. In this experiment, the control laser simultaneously served as the optical pump, resulting in a duty cycle just over 5/6. Our solution was to measure the control laser in a time-resolved way, using a standard photo diode (Thorlabs DET10A, 1 ns rise-time) and sufficiently fast oscilloscope (LeCroy Waverunner HRO 66Zi, 600 MHz bandwidth, 12 bit resolution). The signal was sent to the computer which produced an error signal by integrating manually selected “off” regions, again via Labview, after which everything proceeded as above. This method is cumbersome, provides only slow feedback, and occupies a rather good oscilloscope. At a later time after the hyperfine memory experiments were concluded I tested an alternative method where I switched the EOM to transmit for optical pumping by applying a fast sine-wave to it, switching between it and the RF pulses with an electronic switch (Mini-Circuits ZASWA-2-50DR+). As the low-duty cycle locking method is comparatively slow, and the modulation is symmetrical about the optimal bias, it is possible to maintain a lock in this manner. Unfortunately, the extinction ratio does suffer, and it is a much better solution to optically pump with a dedicated laser.

### 3.2.2 Arbitrary Waveform Generator

A memory experiment performed with attenuated laser pulses is a deterministic experiment in the sense of section 2.2.1. One device serves as master clock, defining the experimental sequence by triggering and synchronizing the others. In the hyperfine memory setup, this master device is an arbitrary waveform generator, the Tektronix AWG7122C. This device is the heart of the experiment’s electronic control which warrants it some description, as well as a few remarks on its principle suitability for memory experiments. The AWG7122C is a massive, pricey, and very general purpose device. With two 8 bit RF channels and four 1 bit markers it can fill the role of two independent waveform or pulse generators and a high timing-resolution digital delay

### 3.2. Pulse Generation

---



generator, all internally synchronized. The sampling rate of  $12 \text{ GS s}^{-1}$  enables the generation of fast, shaped pulses, and the implied shortest width  $\tau = 1/12 \text{ GHz} \approx 83 \text{ ps}$  is well reproduced. In particular, direct electronic measurement using a briefly borrowed sampling scope (Tektronix DSA8300) yields Gaussian pulses with a full width at half maximum of  $86 \text{ ps}$  for a programmed waveform consisting of a well isolated single sample at full output<sup>4</sup> (1 in 60 on, the rest off, repeated continuously). Moreover, it was available on a long-term loan from the quantum dot side of our photon-storage collaboration, allowing us to get started right away and defer the decisions on which specialized devices to buy until we knew what was really needed. Over time we proceeded to learn about the significance of the insertion delay of, jitter to, and maximum rates of asynchronous external triggers, specifically in terms of compatibility to heralded (probabilistic) photon sources – for details see section 5.2.1. In these regards specialized devices perform much better.<sup>5</sup> Further, to reach the required voltage output, this being our EOM’s  $V_\pi$ , external amplification is required. Specialized devices often integrate their amplifiers when the targeted application warrants it, which is not just convenient but increases the reliability of the specifications.

Due to the need for further electronic components before the RF signals reaches the EOMs, in particular we use an amplifier (MiniCircuits ZX60-43-S+) to reach output closer to  $V_\pi$  and a bias tee (ZFBT-4R2GW-FT+) to add the DC offset for locking in each channel, as well as the EOM response itself, we stick to optical characterization of the experimental sequence and pulses detailed below. Nevertheless, this sequence is conducted by the AWG, its marker outputs serving as triggers for further optical switches (acousto-optic modulators), a start signal for the time-to-digital converter, and if need be, gates for the photon detectors. To program the trigger sequence for the markers and pulse shapes for the RF outputs we initially used a large collection of LabView VIs, and I later built a more integrated GUI MatLab app to simplify operation. Either way we can define the output by digitizing various pulse functions, including Gaussian, Lorentzian, exponential, log-normal, quadratic, and rectangular shapes. The widths and amplitudes can be set according to the device’s capabilities. Once measured, the total insertion delay experienced by each channel and marker can be compensated to temporally align the ultimate optical response. Communication over VISA, and saving settings for documentation or later reuse is also implemented.

---

<sup>4</sup>This Gaussian limit of generatable short pulses motivates theoretically optimizing storage for Gaussian control pulses. Other shapes will be well reproduced at ns timescales, but Gaussian pulses will be the most accurate. We are also fortunate that the memory efficiency’s dependence on shape in this sense is weak, an optimal Gaussian is almost as good as a fully optimal control pulse. Critically this also narrows the parameter space for experimental investigation short of implementing a feedback scheme.

<sup>5</sup>While I have not investigated the matter very far there seems to be an underlying trade-off at play here with regards to the device record length (speed vs. space, the old battle first detailed in section 1.1), and thus ultimately the general purpose manufacturers’ main customer demand.




**Figure 3.2:** *top* Conceptual sketch of the hyperfine memory setup displaying the division into fiber-linked components. *bottom* Photo of the setup. Without particular optimization with regards to compact construction, the footprint of the experiment is about  $80\text{ cm} \times 150\text{ cm}$  on the optical table plus a mobile rack mount containing the electronics, including the AWG, oscilloscope, and computer, as well as the detectors in a blackout box.

### 3.3 Optical Setup

The full optical setup of the hyperfine memory experiment is portioned into 7 fiber-linked pieces to improve stability and operability. These are sketched in figure 3.2. A photo of this experiment’s corner of the optical table is also included for a more complete impression. The control and signal laser preparation stages, mounted atop a black breadboard from a time before we had succeeded in claiming this space, dominate its left half. The oven containing the vapor cell at the heart of the experiment can be identified as the prominent cream colored box on the right, surrounded by the polarization optics identified in the sketch. In the upper left background, the control amplification stage (TA) and auxiliary elements can just barely be made out. Next to the table on the upper right of the photo, a portion of the mobile rack housing the electronics

### 3.3. Optical Setup



**Figure 3.3:** Sketch of the signal pulse preparation before the memory and characterization setup.

is visible. Finally, in the lower right foreground, the spectral filtration courtesy of an etalon can be seen. All fibers are single mode and polarization maintaining. The coupling is variably achieved with integrated collimators (Schäfter + Kirchhoff) or ones assembled from components. In the later case a waveplate is required to align the light polarization to the fiber axis. In the sketch, the control/pump and signal boxes represent home built external cavity diode lasers and the division of their light into various fibers as seen on the breadboard. The auxiliary box summarizes a Doppler-free Rb spectroscopy, an offset lock between signal and control, and an additional fiber port for use with a wavelength meter. Connected through the fibers of the signal EOM, the calibration/attenuation stage (not visible in the photo) attenuates the signal with neutral density filters and divides it by polarization for intensity calibration on a single-photon detector. On the control side, the EOM is followed by a tapered amplifier. The TA stage consists of mode-matching optics to improve the amplification and spectral filters to clean amplified spontaneous emission from the output. The basic configuration of the memory itself is sketched in free-space, highlighting the polarization and spatial filtering. The angle between the signal and control beams is  $10(1)$  mrad. An etalon for spectral filtering is again separated out into its own fiber connected stage. The detectors box, this time literally a box to shield the devices from stray light, contains single photon counting silicon avalanche photo diodes, one to monitor the signal and two arranged in Hanbury Brown and Twiss configuration to detect the memory output. The rest of this section divides roughly along those lines.

#### 3.3.1 Signal Pulse Calibration

The setup components to prepare the signal pulses are sketched in figure 3.3. A home built external cavity diode laser<sup>6</sup> (diode: Thorlabs L785P090) using an interference filter (Laseroptik, FWHM  $\approx 400$  pm at 795 nm, angle tunable) as frequency selective

<sup>6</sup>This iteration of the design of these lasers is due to Andreas Jöckel. More information is available online at [atom.physik.unibas.ch/en/people/group-members/andreas-joeckel](http://atom.physik.unibas.ch/en/people/group-members/andreas-joeckel).

element and with an optical isolator (Qioptiq FI-780-5 SV, > 30 dB isolation) integrated in its housing outputs around 35 mW of power on the Rb D<sub>1</sub> line. It is controlled by commercial analog laser controllers (Toptica DCC110 current control, DTC110 temperature control, SC110 piezo scan control, all in Sys DC 110 supply rack). A few milliwatt are immediately fiber coupled for frequency measurement and locking. The rest of the laser power is coupled into the EOM, which can turn the CW output into pulses. After the EOM a beamsampler picks off a small amount of light for locking to the optimal bias point as described above. Absorptive neutral density filters attenuate the pulses to the single photon level. We use a combination of high extinction plates on a flip mount to set the correct order of magnitude and a continuously variable filter wheel (Thorlabs NDC-50C-4) for fine tuning the intensity. A half-wave plate and polarizing beam splitter are used to achieve an equal power ratio after fiber couplings to the memory and the detectors respectively.

For monitoring the photon intensity we use a single-photon avalanche diode or SPAD (Excelitas, SPCM-AQRH-16) with an efficiency at 795 nm specified to  $\eta_{\text{det}} = 60(6)\%$ . The relative uncertainty of the detector efficiency dominates over intensity drifts and determines the accuracy of the calibration. As this detector does not resolve the photon number, the calibration of a coherent state amplitude relies on setting the ratio of the light's vacuum component using Poissonian statistics. The Fock basis representation of a coherent state,

$$|\alpha\rangle = \exp\left(\frac{-|\alpha|^2}{2}\right) \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle, \quad (3.2)$$

immediately yields the probability of having  $n$  photons as,

$$p(n) = |\langle n|\alpha\rangle|^2 = e^{-|\alpha|^2} \frac{|\alpha|^{2n}}{n!}, \quad (3.3)$$


where  $|\alpha|^2 = \langle n \rangle$ . This means that for a state containing one photon on average the probability of detecting nothing is  $p(0) = 1/e \approx 36.8\%$ , while the probability of detecting one or more photons is  $p_{\text{click}} = 1 - p(0) \approx 63.2\%$ . These probabilities are altered when multiple detectors, e.g. two arranged in HBT, are used to calibrate the photon number, as there is then some chance to detect the  $|n > 1\rangle$  components of  $|\alpha\rangle$ . For a pulse repetition rate of  $f_{\text{rep}}$ , the target detection rate to calibrate  $\langle n \rangle = 1$  with a 1 : 1 power ratio in probe and calibration arms is thus

$$c_{\text{det, target}} = f_{\text{rep}} p_{\text{click}} \eta_{\text{det}}. \quad (3.4)$$

Calibration to other mean photon numbers, particularly larger ones, can be performed in the same way, i.e. with  $\langle n \rangle = 1$  in the calibration arm, by changing the power ratio. This insures against decreased accuracy due to detector saturation (too little vacuum

### 3.3. Optical Setup

---

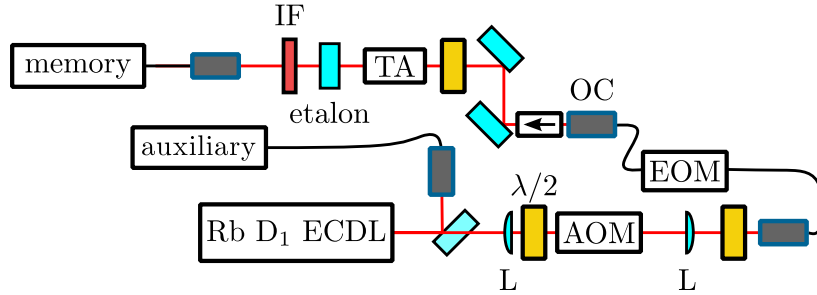


**Figure 3.4:** Time resolved histogram of the signal pulse shape (162 ps bins) and its numerical Fourier transform in terms of the relative power spectral density. As must be expected from EOM generated pulses, a small amount of ringing is visible after the main pulse feature. The curve in the latter is an interpolated spline as a guide to the eye.

component). To avoid including any background caused by the finite extinction ratio of the EOM, only the photons in a time window just large enough to fit the optical pulse are counted. The significance of this background is determined by the ratio of the EOM’s extinction and the optical duty cycle. In our case these figures are comparable which results in a similar number of counts accumulated during the pulse’s off and on times.

The jitter of the SPCM-AQRH-16 detectors is specified to be at least 350 ps. Therefore, to accurately determine the pulse shape a faster SPAD (MPD PDM) with 35 ps jitter is used in its place. This measurement is then limited by the time-to-digital converter’s resolution (Qutools QuTau). While its average bin width is 81 ps, this device has a considerable differential non-linearity. The smallest constant bin size is therefore  $2 \times 81$  ps, combining adjacent internal bins. Counts in histograms using the internal bin width seem to oscillate widely from bin to bin even for constant signals due to this technical limitation. I thus choose to work with the best constant (and actually known) resolution and treat this tagger’s minimum bin size as 162 ps. The measured pulse shape when programming a 1 ns wide exponential decay in the AWG, alongside its numerical Fourier transform is shown in figure 3.4. By interpolation, the 10 : 90 (90 : 10) rise (fall) time is estimated to be 528(40) ps (916(40) ps). The spectral width at the 95 % level, coarsely estimated by interpolation of the Fourier transform, is around 680(20) MHz. These values are influenced by the RF components downstream from the AWG, but are comparable with the envelopes of quantum dot photons as is, and the pulse duration and bandwidth are at least approximately the inverse of each other. To limit the scope of the investigation we do not vary the signal pulses.





**Figure 3.5:** Sketch of the control pulse preparation before the memory. The label L indicates a pair of  $f = 100$  mm achromatic lenses to improve AOM switching performance. The label OC designates anamorphic optics directly after the fiber outcoupling lens to produce an elliptical beam for the TA.

### 3.3.2 Control Pulse Amplification, Optical Pumping

The setup components to prepare the control pulses are sketched in figure 3.5. The components before the EOM are similar to the setup in the signal arm. The laser is the same, except that it is equipped with a higher power diode (Thorlabs LD785-SH300) to produce around 100 mW of power at its output on the Rb  $D_1$  line. Further, an acousto-optic modulator (Gooch and Housego 3080-122, 25 ns rise and fall time) is also used for optical switching to improve the extinction ratio of the control light on the atoms. To achieve the specified speed, the beam is focused into the AOM and collimated after it by a pair of achromatic  $f = 100$  mm lens doublets. The first diffraction order of the AOM ( $\approx 65\%$  efficiency) is then coupled into the EOM fiber. The CW output and hence maximum plausible peak power when pulsing is now typically measured at 8.5 mW. The on-resonance control Rabi frequencies (in units of  $2\pi$ ) we expect to require for optimal storage by modeling fall in the range 1 GHz–2 GHz. These parameters are related by equation 2.4, which implies the required intensity to be achieved by focusing. Unfortunately, at this power the order of magnitude of the required control beam waist to reach such an  $\Omega$  falls into the tens of micrometers. This in turn would mean a Rayleigh range  $< 1$  mm and thus significant spatial variation of the memory performance over the vapor cell. We therefore opt to amplify the pulses with a tapered amplifier (Toptica TA-780-2000-5) so that we can focus less tightly. While specified to 1.5 W CW power at maximum current, multiple issues conspire to reduce the TA’s actual output. First, our working wavelength of 795 nm is far to the edge of its near-Gaussian gain region, which is peaked at 781 nm with a 10 dB-below-peak bandwidth of 16 nm. Secondly, the injected peak power is below the power required to saturate the TA in CW, i.e. “normal operation”.<sup>7</sup> To increase

<sup>7</sup>Amplifying pulses is not the application the manufacturer of this TA had in mind, but has been demonstrated with a model from the same series elsewhere [374], as well as establishing itself as a suitable technique for amplifying considerably shorter ps pulses more generally [375, 376].

### 3.3. Optical Setup

---


the pulse amplitude anamorphic optics (Schäfter + Kirchhoff 5AN-2.5-V-05) are used with the EOM fiber outcoupler to produce a horizontally oriented elliptical beam with form factor 0.4. This improves the performance with this particular gain chip at this wavelength, but does not necessarily translate to other situations. Furthermore, the EOM is protected from the TA's reverse emission by an optical isolator. In hindsight this is overly cautious – an interference filter would be more economical here as reverse emissions of amplified narrowband light are prevented by isolators after the TA.

The TA produces a broad background spectrum of amplified spontaneous emission over its gain region at a level specified to 41 dB below peak when well seeded. As we send pulses into it, pulses which do not seem to saturate the chip even at their peak no less, we expect greater levels of background emissions. To determine their level experimentally a spectrometer with sufficient dynamic range (an optical spectrum analyzer) would be required. The significance of this background is two-fold. On the one hand, light within this broad spectrum that can address the atoms can disturb the atomic storage state. On the other, light at the signal frequency could pass through the spectral filters after the memory and contribute to noise. We therefore frequency filter the output using an interference filter (Laseroptik, as above) removing the broad spectrum and a temperature controlled etalon (2.00(5) mm plate, see table 3.1 in section 3.3.4 for details) to further narrow the spectrum to about 1 GHz. The control light is then fiber coupled to bring it to the memory. The pulse intensity is estimated by the measured maximum CW power of 120 mW, and the components following the TA are optimized in this mode. As this laser is also used for optical pumping it has a high duty cycle during memory operation, which makes this measurement method fairly reliable. I will describe a method better suited for low duty cycle operation due to thermal effects in chapter 5.

Optical pumping to the hyperfine ground state with  $F = 1$  is also performed with the control laser by depleting  $F = 2$ . Before every storage attempt, the control laser is allowed to pump the atoms for 500 ns. The available intensity is sufficient to do so efficiently even though the laser is detuned from  $F = 2 \rightarrow F' = 1$  by  $\Delta$ . To allow for relaxation from the excited state, we leave a 25 ns gap between switching off the pumping and the arrival of the signal.

#### 3.3.3 Vapor Cell and Memory

The atomic part of the memory is quite simple, and all the main components are already shown in figure 3.2. The vapor cell is a 37.5 mm long cylinder with an inner diameter of 5 mm, made by Precision Glassblowing. Two photos are shown in figure 3.6. The windows are flat and anti-reflection coated. Inside is isotopically enriched rubidium with  $> 99\%$   $^{87}\text{Rb}$  specified, as well as 11 Torr of  $\text{N}_2$  buffer gas. The latter is known



**Figure 3.6:** *left* Side and *right* frontal view of the vapor cell mounted in its PEEK oven with partially removed covering. For scale consider that the inner diameter of the cell is 5 mm.

to have a beneficial effect on the optical pumping efficiency by mitigating radiation trapping [377]. By equation 2.8, this buffer pressure increases the time the control laser can address the atoms by about a factor 8, but as we will see in section 3.4 the lifetime will be cut much shorter due to an angle between the spatial modes of the signal and the control. The cell is mounted in a PEEK oven heated with ceramic cartridge heaters (Watlow C1A-9600 20 W) embedded in a copper mount. The temperature is controlled and stabilized with thermistor feedback from the metal with a commercial driver (Meerstetter Engineering, TEC-1091). I will describe a method for determining the atomic temperature directly via spectroscopy in section 5.4 as it was not a part of the characterization of this memory while it was operational. In particular I estimate that the atomic temperature corresponding to the heater temperature of 75 °C used to obtain the main storage results is 63(1) °C. The optical depth on the  $F = 1 \rightarrow F' = 1$  storage transition measured for the experimental conditions after optical pumping is 5.

The signal and control beams are outcoupled from their respective PM fibers in the orthogonal polarizations of the memory scheme and combined on a thin film polarizer (Qioptiq DSPOL VIS-NIR coating). All losses from the signal outcoupler onwards through the memory impact the memory's efficiency, which is particularly significant for single-photon level signals. The focusing of the signal and control beams is therefore performed directly with the fiber outcoupling lenses to minimize the number of optical elements. We perform a spatially resolved measurement of each spatial mode with a beam profiler (DataRay, WinCamD-UCD12, 4.65  $\mu\text{m}$  pixels, relative error dependent on beam size) in an equivalent distance to the vapor cell. The beams are redirected just before the vapor cell with a flip mirror to ensure their foci are in the cell center. The beam diameters in the foci, which are 400(6)  $\mu\text{m}$  for the signal and 525(6)  $\mu\text{m}$  for

### 3.3. Optical Setup

---

the control at the  $e^{-2}$  intensity level, are also determined in this way. Further, an angle of 10(1) mrad between the beams in the horizontal plane is set using this reference. This yields a peak Rabi frequency of the control pulse, by way of equation 2.4 and with a dipole moment of  $d = \frac{1}{2} \times 2.54 \times 10^{-29}$  C m on the addressed transitions, of  $\Omega = 2\pi \times 550(30)$  MHz, assuming a 5% relative error in the measurement of the peak laser power. Reflecting on the experimental results and the simulation output, as well as considering the method of estimating the peak laser power, it is plausible that this estimate is on the lower side of the true Rabi frequency.

After the vapor cell, the lion's share of the control intensity is separated from the signal with another thin film polarizer. The separated control is monitored on a standard photo diode (Thorlabs DET10A), which provides the data to generate an error signal for locking the control EOM. A lens with a long focal length ( $f = 500$  mm) is used to pseudo-collimate the memory output, as some propagation in free-space is desirable to increase the spatial separation between signal and control modes. This has a sufficiently positive effect on the following fiber-coupling leading to the spectral filter given the available in-coupling lenses to justify the additional element here, but should generally be avoided by matching the coupler lens to the mode of the diverging signal beam. For initial temporal pulse alignment and preliminary characterization with weak signals a sensitive Si avalanche detector (MenloSystems APD210), accessed via flip mirror, is used at this point.

#### 3.3.4 Filtration

The control and noise are filtered from the signal by three means: spatially, spectrally, and by polarization. The thin film polarizers used to combine and split the signal and control beams are specified to achieve a 1 :  $10^5$  polarization extinction ratio in transmission. The signal and control polarizations are therefore reversed before the second beam splitter, and any homogeneous rotation or polarizer-axis misalignment is compensated with waveplates. This introduces the main limitation to achieving this extinction, polarization rotation which varies spatially over the extent of the beam diameters. A particularly worrisome source of such variation is strain-induced birefringence from the mounting of the polarizer, which we attempt to address here by attaching the polarizers to their platform mounts gently, with a drop of nail polish. My later comparative tests with even better polarizers favors using a small piece of double-sided tape. Moreover, the waveplates are mounted in precision rotation mounts glued to kinematic mounts for adjusting the angle of incidence. This yields exact retardance control and allows us to reach the specified extinction.

For spectral filtering of noise with a frequency narrowly distinct from the signal, etalons provide a suitable option. As the storage of broadband pulses is under inves-

tigation, the bandwidth passable through such spectral filters is an important figure. Fortunately, polished glass plates of reasonable thicknesses can be made to have suitable properties, specifically high transmission on resonance with bandwidths of 100s of MHz to GHz, with high suppression at the 6.8 GHz hyperfine ground state splitting valued detuning from the signal where the control laser is. Moreover, through thermal expansion the thickness of fused silica can be modified to alter the frequencies resonant to the etalon for light at the Rb D<sub>1</sub> line by 2.391 GHz K<sup>-1</sup>.<sup>8</sup> This enables the design of frequency filtering etalons consisting of single, dielectrically-coated glass plates, tunable through an entire free-spectral range so that the signal may be passed by varying the temperature by just a few tens of degrees. The etalon plates we use are made on demand by the optical workshop Bernhard Halle. They are 15.0(1) mm diameter plates, flat and parallel to  $< \lambda/20$  peak-to-valley at 633 nm, and with the thickness specified either to 0.05 mm or 0.01 mm precision. Dielectric coatings designed for a reflectivity of  $R = 95(1)\%$  at 795 nm are applied to the surfaces. The free spectral range in units of frequency is directly determined by the thickness, with  $\Delta\nu_{\text{FSR}} = \frac{c}{nL}$  where  $L$  is the optical path length, or twice the plate thickness, and  $n$  is the refractive index, with  $n_g = 1.453$  for fused silica at 795 nm [378].

The bandwidths and peak transmissions through these plates, in contrast, are more involved functions of the exact reflectance, material absorption, surface quality, residual power (error in surface curvature), and tilt (error in surface parallelism), and the accuracy of the specifications are insufficient to calculate the typical performance. There is also some amount of variation between plates, not only ones of different thicknesses or different production runs. Across all thicknesses peak transmissions vary between 80% and just over 90%. The bandwidth is measured by scanning a spectroscopy referenced laser over the transmission peak and yields a typical “experimenter’s finesse”, i.e. the ratio of the FSR and bandwidth  $\mathcal{F} = \Delta\nu_{\text{FSR}}/\delta\nu$ , between 45 and 50. I have collected the values for the etalons used in the course of this thesis in table 3.1. Here the thicknesses are specified, the FSRs are calculated from the thicknesses, and the bandwidths are measured as I just described, excepting the 1 mm etalon where I have not measured anything but obtusely followed the doubling trend in the table nonetheless. The suppression is measured in CW and at low power, around 1 mW. At high laser power the performance is generally somewhat worse due to the variable thermal load.

To temperature stabilize and tune these etalons we use Peltier modules with central boring for the beam (UWE electronic UEPT-440-127-079E120) connecting an aluminum “object” holder, with the glass plate locked inside by retaining ring, to a


---

<sup>8</sup>As thermal expansion, i.e. a change in resonator length, is at play this figure depends both on the etalon material and the wavelength of the light. For example, for light at the Rb D<sub>2</sub> line our etalons temperature tune by 2.437 GHz K<sup>-1</sup> instead.

### 3.3. Optical Setup

Thickness, mm	FSR, GHz	Bandwidth, MHz	Suppression at 6.8 GHz, dB
8.00(5)	12.8(1)	280(10)	-32
4.00(5)	25.5(4)	550(10)	-26
2.00(5)	51(1)	1150(20)	-18
1.00(5)	102(5)	about 2300	-10

**Table 3.1:** Table of etalon parameters relevant to spectral filtering.



**Figure 3.7:** *left* Original and *right* most recent design iteration of our etalons. The cumulative upgrades include embedding the thermistors inside the holders for greater measurement stability, brass retaining rings to extend the expected lifetime of the threading, copper sinks for a slightly larger stable tuning range, 3D-printed mounts for the TEC controllers for thermal insulation and to mitigate the risk of damaging them while mounting, and anodization of the aluminum object holder to prevent stray reflections.

larger “sink” block of aluminum or copper. Copper sinks allow for stable operation in a slightly larger temperature range, which is useful for thin plates. These metal parts are machined for us by the in-house mechanical workshop, and the design has gone through a series of iterations. Some photos are shown in figure 3.7. The sink and object temperatures are measured either on the front side of the mount or in small holes drilled into the metal. We use PT100 standard thermistors in 4-wire configuration for the object temperature measurement, and NTCs for the sinks as this measurement is exclusively needed to switch the TEC off should the sink get too hot. The mechanical connections are made with a heat-conducting two-component epoxy. As with the vapor cell oven, the temperature is controlled and stabilized with a commercial driver (Meerstetter Engineering, TEC-1091). The measured temperature fluctuations typically remain  $< 10$  mK up to around  $60^\circ\text{C}$  and  $< 5$  mK near room-temperature. This is a testament to the driver capabilities as these mounts are completely exposed to the air. The temperature range required to tune one free spectral range is just  $\Delta\nu_{\text{FSR}}/2.391 \text{ GHz K}^{-1}$ , meaning that the thickest plates always can be held within

3°C of room-temperature. The frequency stability implied by the thermal expansion coefficient compares favorably with the bandwidth, fluctuations  $< \delta\nu/20$  are expected in every case.

Now for some practical remarks on achieving the stated performances. To improve the mode overlap of the beam with itself within the etalon for optimal performance it pays to use a minimally divergent spatial mode. A collimated beam works well, but a so called pilot beam, which is a beam focused to have its waist at half the working distance, will have the most constant beam diameter over that working distance. When placing the etalon after a fiber outcoupler and beginning with a collimated beam, slight adjustments to the outcoupling lens after other alignment options are exhausted can improve the transmission by another few percent. Finally on temperature, it is important to account for the thermal load on the etalon induced by high laser power. In particular aligning an etalon to remove the broad spectral background from a CW TA at full output power does not result in good transmission of low duty cycle pulses without a significant temperature readjustment. For tips on aligning many etalons in series see section 5.2.5.

The hyperfine memory uses just one 4 mm etalon for output filtration. It is reached via a fiber after the memory, and after the etalon another fiber leads the signal to the detector. As indicated in the table, this filter does not provide a large amount of control suppression. Indeed its function is dual, as temperature scanning it during memory operation yields some coarse spectral information about the output. The total control suppression measured at the output of the fiber leading to the detector is at least 120 dB, the remaining extinction stemming from the 10(1) mrad angle between the spatial modes of signal and control. The effect of this angle is discussed in section 3.4. About 2/3 of the signal is also lost in transmission through the memory (5 dB loss), losses dividing roughly equally between the etalon, the two fiber-couplings, and all other reflective and absorptive losses. To ensure an accurate measurement of the technical loss, excluding the atoms, this figure is measured in CW with a probe power of a few mW with the vapor cell held at room temperature. It is thus a lower bound on the loss the weak coherent pulses are subject to.

### 3.3.5 Auxiliary Setup Elements

The memory operation is sensitive to the two-photon resonance condition. We thus ensure that this condition is met by offset locking the signal and control at the hyperfine ground state splitting.<sup>9</sup> Light from the auxiliary ports of the signal and control lasers

---

<sup>9</sup>A skeptical observer might now wonder whether this kind of frequency lock between the control laser and an actual single photon source is realistic. In practice a lock would have to be implemented between the control and the device responsible for the generation of single photons, i.e. the excitation laser for quantum dots or the pumping laser for SPDC pair sources. In the later case this would imply

is combined into one spatial mode on a 50 : 50 beam splitter to generate a beat note. This beating is detected with a fast photo diode (Hamamatsu, PD G4176-03), which is manually biased via bias tee (MiniCircuits, ZX85-12G-S+). The diode signal is down-mixed (MiniCircuits ZMX-7GMH) with a local oscillator (CTI, PDR0-6378) specified to 6835(5) MHz, and the error is measured with a frequency counter (aim-TTi, TF930 3 GHz). Feedback to the control laser piezo is generated by the computer integrated NI-card. This fixes the relative frequency of the signal and control lasers. The absolute frequencies we want to investigate span a range of a few gigahertz around the line, and the memory performance depends only very weakly on the detuning at the level of the short term laser drift. We therefore tune the working point  $\Delta$  by referencing a wavelength meter (HighFinesse WS-7) via another fiber connection in the auxiliary setup and leave the signal laser free-running. As an additional reference we also implement a Doppler-free Rb spectroscopy for both lasers here, used for instance to calibrate the etalon bandwidth measurements.

### 3.4 Results

In this section I will present the main results we reported on in [365], the figures of merit of this hyperfine memory. The parameters used to obtain those results were chosen on the basis of some empirical optimization. I will defer discussing this optimization and characterizations to support our conclusions about the limitations to section 3.5, so as to first convey how the data are obtained and analyzed in a specific example.

The experiment is performed at a detuning  $\Delta = -2\pi \times 900$  MHz from  $F' = 1$ . The experimental sequence is repeated at a rate of  $f_{\text{rep}} = (600 \text{ ns})^{-1} = 1.6$  MHz. It begins with 500 ns of optical pumping by the detuned control, followed by 25 ns of waiting for the excited state to decay. Then an attempt is made to store an approximately 1 ns wide,  $\langle n \rangle = 1.0(1)$ , exponentially-decaying signal (shown in figure 3.4) with a 5 ns FWHM Gaussian control pulse, for 50 ns. Finally a second, identically-generated control pulse reads out the signal and the sequence restarts after about another 25 ns.

The results of this process are measured by two single-photon avalanche diodes, or SPADs, (Excelitas, SPCM-AQRH-16,  $\eta_{\text{det}} = 60(6)\%$ ) in HBT configuration. This detector arrangement allows us to measure the second-order coherence. However we are also interested in arrival-time histograms, both for direct tuning feedback and as

---

a transfer cavity scheme as an SPDC pumping laser must have a considerably higher frequency to produce photons on the Rb  $D_1$  line. Fortunately, given some warm-up time, our lasers drift slowly, typically  $\leq 10$  MHz in 10 min. In particular this is slow in comparison to the required integration times for statistically significant data. I abandoned this locking scheme a short time after the hyperfine memory measurements were published and found no practical reduction in performance, concluding that such a lock might not be required at all given sufficient inherent stability of the single photon source, at least not for characterizing measurements.



datasets from which to evaluate the memory efficiency (given by equation 1.9) and lifetime therein. As these detectors are non-number resolving, the detection probability of small amounts of coherent light such as our signals are altered by the detector pair arrangement from what they would be on a single SPAD with the efficiency  $\eta_{\text{det}}$ . A quick calculation is warranted. Consider a weak field with amplitude  $|\alpha|$ . A 50 : 50 beam splitter is a device that splits the intensity in half, the field amplitude in each output is therefore  $|\alpha|/\sqrt{2}$ . Far from saturation, the detector efficiency can formally be absorbed into the field amplitude as well to yield a detected amplitude of  $\alpha' = \alpha\sqrt{\eta_{\text{det}}/2}$ . Now, as the detectors are non-number-resolving, we must work backwards from the vacuum component like in section 3.3.1. Using equation 3.3 the probability of detecting no photons in a coherent state is  $p(n = 0) = \exp(-|\alpha|^2)$ . Therefore, if the photon arrival-time histograms of two detectors in HBT are summed, we should expect to detect

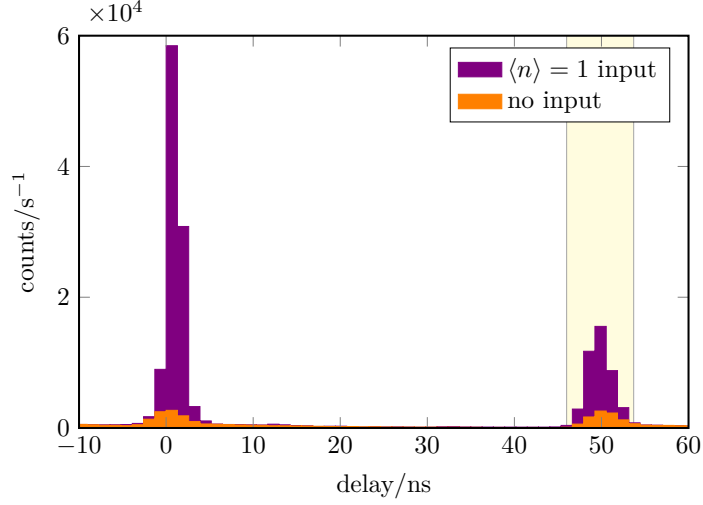
$$n_{\text{det}}^{\text{HBT}}(|\alpha|^2) = 2 \left(1 - e^{-|\alpha'|^2}\right) \quad (3.5)$$

photons per pulse. In particular, this yields  $n_{\text{det}}^{\text{HBT}}(|\alpha|^2 = 1) = 52(6)\%$  for  $\eta_{\text{det}} = 60(6)\%$ , which is noticeably higher than the corresponding  $p_{\text{click}} \times \eta_{\text{det}} = 40(4)\%$  for a single detector. In the terms of equation 1.9 this value takes the role of  $\eta_{\text{source}} \times \eta_{\text{det}}$ .<sup>10</sup> The other parameters,  $N_{\text{in}}$ ,  $N_{\text{out}}$ , and  $N_{\text{noise}}$  are all directly measured numbers of counts (equivalently, rates may be used if the integration times are not exactly equal). To reduce the data volume, the periodicity of the experiment can be exploited. A sync divider is employed to register only 1/128 of the triggers representing the start of the experiment, from which the complete set can be reconstructed. The repetition rate is sufficiently stable that potential timing errors induced by this are negligible in comparison to the detector jitter. It is therefore convenient to write  $N_{\text{in}} = f_{\text{rep}}t_{\text{int}}$ .

Figure 3.8 shows a typical histogram obtained from the described experiment, restricted in scope to show the photon counts retrieved after 50 ns and those leaking through the memory as storage is attempted at  $\tau = 0$ . As both the timing and temporal width of the output is known, the photons are counted in a selected region of interest (ROI). This ROI is also called a coincidence window in analogy to a heralded experiment where the signal has a fixed delay relative to the detection of the herald. If necessary, in a real application, either the input of the device receiving the memory output could be gated, or the memory could be switched optically to isolate out these desirable photons. For the noise measurement a beam block is placed in the path of the signal input before the memory. Data are integrated for 10 min, yielding about  $26.5 \times 10^6$

---

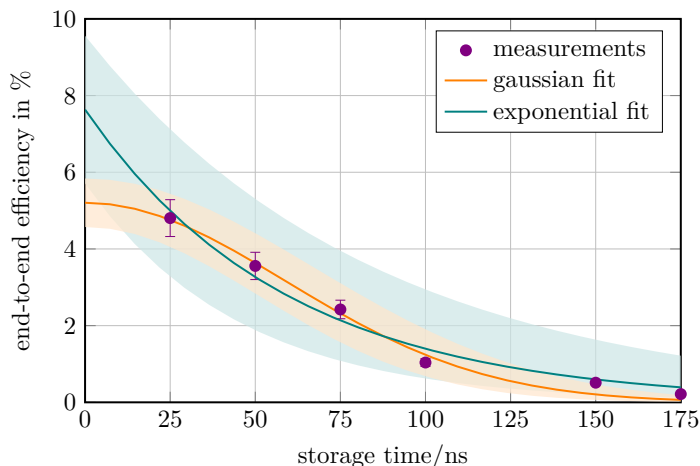
<sup>10</sup>In [365] we simplify the analysis for the sake of presentation by simply treating the  $|\alpha|^2 = 1$  state as an ideal source. That approach neglects both vacuum and detectable multiphoton components, introducing two errors of, roughly, similar magnitudes and opposite directions. The efficiencies reported there are slightly lower than what results from an exact accounting.



**Figure 3.8:** Storage and retrieval of weak coherent pulses with 50 ns storage time. The bin size is 1.3 ns ( $16 \times 81$  ps). The 8 ns ROI for retrieval is shaded yellow. It visibly excludes fluorescence occurring after the signal has been read out. A small amount of unintentional retrieval during the storage time, distinct from read-in induced fluorescence visible in both curves, can be seen in the storage data. The counts are normalized into rates by their integration times, about 10 min each.

counts in the ROI for the storage experiment and around  $5.6 \times 10^6$  counts in the same region and time frame when the input is blocked, which is sufficient for statistically meaningful analysis. The main features of the leaked peak, retrieved peak, and noise peaks are reasonably well fit by Gaussian functions of around 2.5(1) ns, 4.4(4) ns, and 5.7(6) ns FWHM respectively, and are followed by exponential tails with fit decay times around 19(1) ns (errors are 95 % confidence intervals of fits). This exponential decay time is underestimated as its signal fades beneath the background, and would most likely correspond to the 28 ns excited state lifetime if it were not so. Note that the detection scheme is not optimized for low jitter, which leads to systematic widening and the Gaussian shapes.

The memory efficiency evaluated from the counted photons is  $\eta_{e2e} = 3.9(4)$  %. This end-to-end figure is not corrected for technical loss or memory lifetime and requires only the assumptions of accurate input calibration and detection probability. If we want to compare this result to a theoretical model of the memory interaction we must make these adjustments. The signal to noise ratio can also be read off this histogram, it is  $\text{SNR} = 3.7(5)$ . To abstract this value for comparative purposes the source efficiency could again be divided out, however the detection efficiency of the noise would also need to be considered which in turn depends on the nature of the noise. While these histograms were recorded on slower SPADs than when I first presented this method in section 3.3.1, the numerical Fourier transform of the full resolution (162 ps bins)



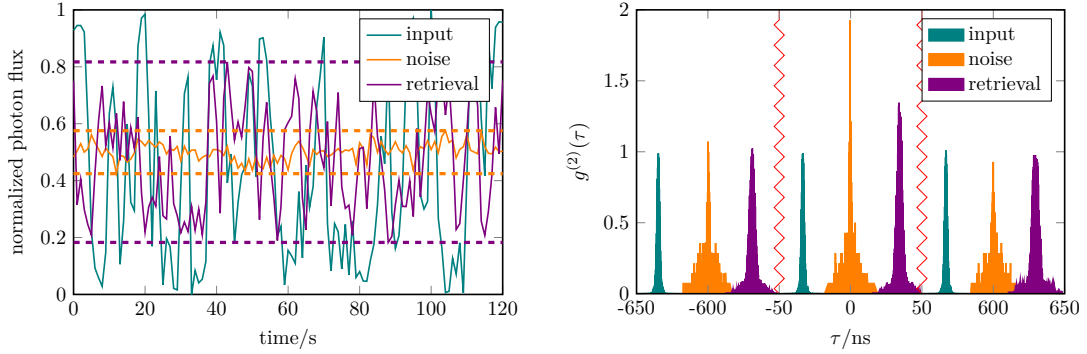
**Figure 3.9:** Memory efficiency evaluated as above for various storage times to measure the lifetime. The Gaussian fit yields a  $1/e$  lifetime of 83(11) ns and an initial efficiency of 5.2(6) %. An exponential function, which fits considerably worse, yields a 60(20) ns lifetime and 7.6(19) % initial efficiency. The shaded regions represent the 95 % confidence intervals of the fits.

pulse shapes can still yield approximate information about the memory’s effect on the input’s frequency domain characteristics. By this method, the bandwidth of the pulse leaked through the memory is about 450 MHz and the bandwidth of the retrieved pulse is about 250 MHz. Due to jitter, both of these values are necessarily underestimations. Nevertheless, this is at least some evidence that the performance here is limited by the control laser’s ability to store the entire input spectrum, and that higher  $\Omega$  or a slightly narrower signal could have significantly improved performance.

The memory is operated with an angle between the signal and the control, and so it is to be expected that the lifetime is limited by spin wave dephasing. With the measured angle of  $\alpha = 10(1)$  mrad and equation 2.5 the expected ( $1/e$ ) lifetime is 70(8) ns. As the physical process at work is motional spin-wave dephasing a Gaussian decay is expected – compare [175]. Figure 3.9 summarizes a series of experiments at various storage times evaluated for memory efficiency like the example of figure 3.8. To minimize error, in particular due to longer-term alignment variations from mechanical relaxation in the optical mounts, lifetime measurements are best performed in one session. A Gaussian fit to the data converges well and agrees with the predicted lifetime. Correcting the efficiency extrapolated to zero storage time for 5 dB technical loss yields an internal efficiency around  $\eta_{\text{int}} = 17(3)$  %. I will compare this performance to a model in section 3.5.


Next we can investigate whether the memory preserves the coherence properties of the input light. The first-order coherence is measured by inserting an unbalanced, fiber-

### 3.4. Results



**Figure 3.10:** *left* Interference between subsequently produced input pulses, read-out noise, and retrieved pulses. Each point represents 1 s integration time. Full contrast is defined by what is achieved for the input pulses. The visibility observed for the retrieved pulses is then  $V = 0.65(5)$ , while that of the noise, presumably due to random intensity fluctuations, remains below 15%. *right* Second order autocorrelation  $g^{(2)}(\tau)$  of photons detected in the input pulse, noise, and retrieved signal, respectively. The read-out signal reaches  $g^{(2)}(0) = 1.35(7)$ . The read-out noise exhibits  $g^{(2)}(0) = 1.9(2)$  and the input yields  $g^{(2)}(0) = 1.01(1)$ , consistent with thermal and coherent light, respectively. Errors are statistical of the type  $1/\sqrt{N}$ . Data are normalized to the peaks at  $\pm 600$  ns and shifted for better visibility, i.e. all measurements actually have their peaks at 0 and  $\pm 600$  ns.

based Mach-Zender interferometer with a 400 ns (80 m of fiber) arm length difference and adjusting the repetition rate to 1/400 ns. This lets subsequently stored pulses interfere and allows us to observe the contrast by monitoring one output of the interferometer as its phase undergoes random thermal fluctuation. This measurement is presented on the left side of figure 3.10. Thereby full contrast is normalized to what is observed when strong laser pulses are sent through the interferometer directly to compensate for imperfect mode matching. The retrieved pulses show a reduced contrast, while the noise shows very little. We can observe that the reduction in visibility of the output scales with the signal to noise ratio, which we vary by increasing the average photon number in the input pulse to the memory. At  $|\alpha|^2 = 10$  almost the full visibility is recovered by the retrieved pulses. Additionally, the data collected in HBT can be analyzed to yield the second-order autocorrelation of the detected light. The results of this measurement are shown on the right side of figure 3.10. This measurement confirms the thermal nature of the noise when the signal is blocked. This verifies that the memory is not limited by leaked control light at least, however beyond that the result is somewhat worrisome. First we must be wary, as the  $g^{(2)}(0)$  of the input is not almost zero, the shorthand estimate of equation 2.6 is not accurate. Thus I turn instead to the incoherent noise model of equation 2.7. This, however, yields an expected value of only  $g^{(2)}(0)_{\text{ret, theo}} = 1.045$ , suggesting that some amount of the noise is produced in




**Figure 3.11:** *left* Internal memory efficiency as a function of the control pulse width at constant peak power. The line is a  $1 - \exp$  shaped guide to the eye. *center* The amount of noise scales linearly with the control power for these parameters indicating fluorescence noise. *right* The temperature of the filter etalon is scanned to provide some spectral information on the noise. The dashed lines indicate the frequencies of (left to right)  $F' = 2$ ,  $F' = 1$ , and the control laser assuming a frequency shift due to thermal expansion of  $-2.391 \text{ GHz K}^{-1}$ . The total frequency range of the scan is 5 GHz. The spectrum is broadened by the large etalon bandwidth of about 1.1 GHz

a coherent process. Models for the retrieved  $g^{(2)}(0)$  for such cases were investigated by others [187], but we did not pursue the matter of modeling further. The failure of the incoherent model is compatible with boogeyman processes like four-wave mixing gain.

### 3.5 Performance Optimization and Noise Analysis

To provide a sample of the optimization procedure for the memory over the large conceivable parameter space I show a few characterizing measurements performed at  $\Delta = -2\pi \times 300 \text{ MHz}$  in figure 3.11. These measurements reasonably represent a number of recurring conclusions. Given that the control pulse duration is sufficient for adiabatic storage, the efficiency becomes a weak function of the pulse width. The transition point depends on the signal and is fairly robust when varying other memory parameters. The minimum width producing a high efficiency is a good choice to minimize the noise, so that the optimization can be performed in a nearly unidirectional manner. This near resonance with the line, the read-out noise grows linearly with the peak control power, indicating that its origin is fluorescence. Scattering of the control light into the signal channel would also show this behavior, however the filtration including the spatial separation of the signal and control modes is sufficient to eliminate this explanation. Although this is weaker evidence than the second order autocorrelation, it matches our expectations for small  $\Delta$ . If sufficient control power is available, it is possible that the best signal to noise ratio is achieved with slightly lower  $\Omega$  than is fully optimal for storage. This is because the growth in efficiency saturates with additional



**Figure 3.12:** Data points: internal efficiency for storage and subsequent retrieval of the experimentally realized signal as a function of the detuning  $\Delta$ . Note that the measured efficiency at  $\Delta = -2\pi \times 900$  MHz was slightly lower in this scan than in the main result. Lines: Simulation with same OD and level scheme as in the experiment for the experimentally realized Gaussian control pulses with 120 mW peak power (violet), Gaussian control pulses with higher laser power, about  $4 \times \Omega_{\text{exp}}$  (orange), and fully optimal control pulses of arbitrary shape (teal). The shaded regions represent the absolute numerical accuracy of the simulation, which is estimated at 1%, with an additional 5% relative error from the measured parameters added to the simulation of the experimental situation.

power near the maximum, while the noise scales as it does elsewhere. At the working point  $\Delta = -2\pi \times 900$  MHz used above, this point is moot as the available control Rabi frequency is considerably lower than required for optimal storage. Finally, we can see that in this configuration, little spectral information is gained by scanning the filter etalon. For this measurement we sought to be quantitative and allowed for plenty of thermalization time for each data point. The result is disappointing. Qualitatively, moderately fast scans (e.g.  $0.05^\circ\text{C s}^{-1}$ ) monitored continuously are more informative. Stokes, anti-Stokes, and much broader fluorescence peaks can be identified in this manner. I will show more informative measurements of this kind in chapter 5 when discussing the characterization of the Zeeman memory, for which I refined the technique to yield more quantitative data.

The most significant optimization we performed on this memory is in the detuning from the excited state. This is shown in figure 3.12 along side numerical simulations for both the experimental situation and two alternative control scenarios. As control

and signal frequencies are offset locked,<sup>11</sup> and as only one etalon spectrally filters the control before the memory, collecting these data is relatively simple in comparison to later setups. The observed variation justifies even considerable efforts to find the best  $\Delta$ . The simulations include four levels and are initially based on work by Matt Rakher in cold atoms [120] implementing Gorshkov’s universal approach to lambda memories in free space [117]. They were modified to the case of hot atoms by Janik Wolters following [118] for [365] and include 500 MHz of Doppler broadening. Moreover the purely absorbing transition  $|F = 1, m_F = 0\rangle \rightarrow |F' = 2, m'_F = 0\rangle$  (see figure 7.1 for details) is included as a parasitic single-photon transition, and its deleterious effect on storage efficiency in the vicinity of  $F' = 2$  is well reproduced. The peak value of the achievable efficiency is also accurately captured, although its exact detuning and the width of the features are off. The simulations with optimal control power support the conclusions about the limitation posed by  $\Omega$ . Simultaneously, only minor improvement is predicted in additionally optimizing control shape, which justifies the choice to use Gaussian control pulses.

On the accuracy of the x-axis in figure 3.12, while my involvement in implementing this simulation was negligible, I did proofread the code and I am convinced that this is not an artifact of an isolated scaling error in  $\Delta$ ,  $\Omega$ , the optical depth, or any of the other parameters that can have variant definitions for experimentalists and theorists along the lines that I pointed out a few times in section 2.1.4. A working theory is that spatial variation of the control beam, i.e. its intensity profile, could be the origin of the discrepancy, as the model is one dimensional. In simulations of EIT in cold atoms, this same problem has been successfully addressed by dividing the medium into ring shaped sections of more constant intensity [379], and fully three dimensional theoretical models exist as well [380, 381]. A new simulation taking radial spatial variation into account was created by Roberto Mottola to model the results of [382] and is reported on in full in [383]. In light of later insights, I also refrain from reproducing the curves in figure 4 of [365] that show the model without single-photon transitions, as they do not accurately reflect a physical scenario capable of this suppression. In particular, any such scenario must reduce the addressable Zeeman states of the  $D_1$  line, which needs to be accounted for in the transitions’ dipole moments. The optical depths producing high efficiencies then require higher atomic temperatures. As we did not yet have a specific scheme in mind at the time this adjustment was not done.

---

<sup>11</sup>In hindsight, the offset lock may raise a concern about the possibility of control induced level shifts – does the fixed frequency difference truly ensure that the double resonance condition is met? During this optimization we did, in fact, attempt to verify that the set frequency difference between signal and control produced the best results. Note, however, that the maximal induced level shifts are only on the order of a few 10s of megahertz, that each value of the two-photon detuning achieves its optimal performance at a different  $\Omega$ , and that variation near the local optima is only slight. Unsurprisingly, then, varying this value slightly did not produce stark changes, and we thus opted to leave it alone.

### 3.6 Towards Storing Single Photons

We started this work, and I started this chapter, envisioning an interface between an atomic vapor memory and a quantum dot photon source. Having detailed the measurements, we can now state the required source efficiency to achieve the bare minimum of a sensible interface, a signal to noise ratio exceeding unity:

$$\eta_{\text{source}} \stackrel{!}{\geq} \frac{p_{\text{click}}}{\text{SNR}_{|\alpha|^2=1}}. \quad (3.6)$$

The correction by  $p_{\text{click}}$  is conservative due to the detection in HBT, so this value should truly set a lower bound. Plugging in the measured values yields the requirement  $\eta_{\text{source}} \geq 17\%$ . This is demanding for a quantum dot source, with typical extraction efficiencies of  $< 1\%$  from the bare heterostructure. A common and relatively simple way to improve this efficiency by as much as an order of magnitude is to use a solid immersion lens, but even an upgrade like that would not yet suffice for interfacing. In section 2.2.3 I mention more involved optical engineering methods for further increasing quantum dot extraction efficiencies, particularly embedding them in cavities. With these techniques the efficiencies could conceivably be made high enough for direct compatibility. However some of these developments are more recent than when we finished our characterization of the hyperfine memory, and the highest values have generally been reached in the “workhorse” system of strain-grown quantum dots emitting around 950 nm. Rubidium-like quantum dots may well be better behaved systems in the grander view, but they have not been made and experimented with for as long, so not every demonstrated technique has been transferred yet. We therefore concluded that significant improvements on both sides were required first, and decided to bridge the gap with a tailored cavity-SPDC pair source primarily optimized for efficiency, while also remaining in the same regime of speed and bandwidth. This could serve both as an interesting interface in its own right, and as a better comparative stand-in than laser pulses for other photon sources that had not been tailored to the memory.

As for the required improvements to the memory, our results laid bare many areas for improvement. While the distinction between possible noise sources was not fully made, we found some evidence of both fluorescence and four-wave mixing. Moreover, at higher optical depths, the latter becomes an ever greater concern. The availability of states that can absorb the signal without storing it is also problematic. While not a main limitation to the observed performance at  $\Delta = -2\pi \times 900$  MHz, it narrows the possible range of detuning and further complicates the use of the optical depth as a tuning knob for the efficiency. Both of these problems can be addressed simultaneously by controlling the Zeeman state of the atoms and exploiting polarization selection rules



[384], so we devised to achieve this control and implement a memory scheme less subject to these flaws. This begins the list of desirable technical improvements. To pump atoms to a Zeeman state in the Rb hyperfine ground state manifold two pumping lasers are definitely required, as both ground states will need to be addressed. Conventionally the laser that is tasked to address atoms in the correct  $F$  but incorrect  $m_F$  state is dubbed the pumping laser, while the laser tasked to address atoms in the incorrect  $F$  state is called the repumping laser. Both of these tasks must be fulfilled, and for the purpose of independent optimization preferably neither should be performed by the control laser. Additional control power is also necessary for maximally efficient operation. The spatial filtering, while principally effective, puts a very low bound on the memory lifetime. Additional spectral filters should be just as able to separate signal and control without introducing this limitation. All of these changes and improvements, as well as many more of which the necessity was not yet apparent here, are the subject of chapter 5.



## Chapter 4

# SPDC Photon Pair Source

*If your experiment needs statistics, you ought to have done a better experiment.*

Ernest Rutherford (apocryphal)

---

Taking stock together with our colleagues working on Rb-like quantum dots in Richard Warburton’s group in 2017, after the experiments with the hyperfine memory, we jointly concluded that interfaceability between our systems was still a ways off. A photon pair source based on spontaneous parametric down-conversion, tailored for compatibility with a memory in hot Rb vapor, seemed like an ideal alternative for the near term. Such an interface would be interesting in its own right, as well as serving as a realistic test bed for interfacing other photon sources. Moreover, striking up a collaboration with Oliver Benson’s group at the Humboldt-Universität zu Berlin who brought a great deal of specific expertise on engineering the properties of SPDC sources to the table, we felt confident that we could achieve compatibility to the demonstrated memory figures of merit with such a source. The original design and construction of this source was led by Janik Wolters, embedding himself in Oliver Benson’s group for a time, while I continued work on the memory side. Mid-September 2018 the first version of the SPDC source, mounted on a thick optical breadboard, traveled some 870 km from Berlin to Basel.

While I had already begun work on implementing the Zeeman memory scheme at the time, I had not yet been able to make its more involved setup reactive to a heralded photon source. These issues and their solutions are discussed fully in sections [5.2.1](#) and [5.2.3](#). Driven by optimism induced by the high source efficiency, I reverted the memory setup to the hyperfine scheme and began interfacing attempts. In this time Roberto Mottola took charge of operating the source, and eventually rebuilt it almost in its entirety due to the non-linear crystal at its heart breaking. By spring 2019 we had stored photons that just barely exhibited the desirable number statistics of

$g^{(2)} < 1$  upon retrieval, but this did not really live up to our hopes, as well as now being even more obviously incompatible with an incoherent noise model. The prospect of first reporting on this source in the context of an interfacing experiment crept farther and farther into the future. Roberto and I therefore proceeded to, for the first time so far, perform a thorough performance characterization of the source output on its own. We reported on this and the work leading up to it in [385]. My involvement lay predominantly in collecting and modeling the statistical results, so I will attempt to be brief on other aspects of this source here.

In what I assume will remain a famously widespread turn of bad luck at the very end of 2019 well into the future, just as I was about to resume interfacing experiments with a memory that still used the hyperfine scheme but resembled what I describe in chapter 5 in most other ways, the source crystal broke once more. Our one remaining replacement, rapidly installed by Roberto, was considerably shorter than the one of [385]. This means it produced photons with nearly 1 GHz bandwidth. While this had little impact on the memory efficiency, it required using correspondingly broadband spectral filters after the memory. This proved to be disastrous for the signal-to-noise ratio – I was never able to improve it beyond one in this system. Apparently the noise was considerably broader than the signals we had stored so far, and I had somewhat unwittingly been removing much of it simply by matching the spectral filters to the signal bandwidth.<sup>1</sup> As a result I finally lay the hyperfine memory scheme to rest, and Roberto and I began thinking about an iteration of the source design that would be more easily operable and less susceptible to damage. Almost immediately Björn Cotting began implementing these ideas during his master’s thesis work, Roberto and me guiding him with the experience we had just recently obtained in the production of [385]. The resulting thesis [386] contains an excellent description and exhaustive characterization of the source almost exactly as it is used in the interfacing experiments I present in the next chapter, and is freely available. I will therefore limit my own discussion of the improved source design to quoting his relevant results and describing a handful of necessary changes for interfacing with a memory. Finally, I will comment on the tunability of this double resonant source to the Rb D<sub>2</sub> line.

## 4.1 Source Design and Expected Performance

As I briefly described in section 2.2.4, the photon generation process in spontaneous parametric down-conversion relies on a medium with a  $\chi^{(2)}$  non-linearity. Therein high frequency pump photons are spontaneously split into pairs of lower frequency photons

---

<sup>1</sup>Of course I knew the noise was broad, but I did not anticipate the effect would be this large. The motivation for matching the spectral filters to the signal bandwidth lay more in the greater suppression thicker (i.e. lower bandwidth) etalons provide at the control frequency than anything else.

labeled signal and idler. The demands set by the memories I cover in this thesis concern four properties: photon bandwidth, signal photon frequency, heralding efficiency, and heralding rate. The photon bandwidth should lie in the 100s of MHz to exploit, but not exceed, the memories' capabilities. This requires having the photon generation process occur in a cavity resonant to the signal. SPDC in cavities is treatable with the theory of optical parametric oscillators (OPOs) pumped far below threshold [387, 284]. Strictly speaking only the signal requires a cavity. Signal-idler frequency entanglement ensures the bandwidth of the photons produced is narrowed even without cavity resonance for the idler [388]. A cavity simultaneously resonant with the idler is useful, however, as it yields Purcell-like enhancement of the photon pair generation rate – an effect which does not occur in single resonance schemes [389]. The signal photon frequency must be at the Rb  $D_1$  line. Tunability within at least 1 GHz of the line is a necessity for meaningful flexibility with regards to finding an optimal detuning from the excited states of Rb, and a few GHz range would be decidedly better so that both red and blue detunings can be investigated. The stability of the signal frequency is also relevant to maintaining the two-photon resonance condition with the control. It determines whether interfaced operation requires a source-memory frequency locking scheme. If the frequency stability is similar to that of a laser, operation on the timescales required for characterization would be possible without a lock, which would be convenient.

An implication of the demand for a single signal frequency is that the source should emit photons in a single mode. More precisely, irrespective of what the source emits, the heralded photons must be of the signal frequency and only of the signal frequency. The practical conditions of this requirement being fulfilled in SPDC sources with cavities for their signal and idler photons are treated in [390], and additional theory on multimode emissions can be found in [391]. In a nutshell, the conservation law based conditions of the down-conversion process (see section 2.2.4) are met in *clusters*, where the cavity FSRs of the signal and idler coincide, as well as weakly where the matching is “only off by one”. To ensure that all heralded photons are in the desired mode, the idler photon can be spectrally filtered. It is shown in [390] that the expected number of clusters over the full bandwidth that enables phase matching is 3, and that if the dominant one is well centered the contribution to the pair generation of the 2 side clusters is only 40%. This filtering is therefore not expected to have too large an impact on the “bare” heralding rate. The signal arm will contain a background of uncorrelated photons for which the corresponding idler was lost or filtered. Interfaced with the memory, the included spectral filters remove the uncorrelated background from additional unheralded modes, but not that present in the signal mode itself. The exact frequency of the idler photon is not important, in particular there is no reason for it to also coincide with the Rb  $D_1$  line. However, something in the near-infrared close to the Rb lines is preferable so that mostly “lab standard” optical elements can be

used. The detection efficiency of the readily available single photon detectors (Excelitas SPCM-AQRH) is also a consideration here, as it falls off more rapidly with increasing wavelength than, for example, the effectiveness of anti-reflection coatings.

Conventionally an external cavity is at odds with the need for a high heralding efficiency. Every surface, including those of external mirrors, produces an opportunity for photon loss and wavefront distortion. While I estimated that a source efficiency of 17% should yield a signal-to-noise ratio of one in an experiment interfacing a photon source and the hyperfine memory at the end of chapter 3, that performance probably would not suffice to demonstrate retrieval of non-classical light. It is hard to overstate how desirable high efficiency is in a photon source for interfacing it with a memory producing some read-out noise. An alternative to external cavities is to make the non-linear medium itself the cavity by polishing and coating its end facets. Monolithic schemes like this had already been successfully implemented at the time of designing the source described here, for instance in [390], although the reflectivities of the coatings therein had not been sufficiently imbalanced to strongly favor pair emission in one direction. Finally there is also the matter of rates. The hyperfine memory was operated at  $1/600 \text{ ns} \approx 1.6 \text{ MHz}$ . To make use of this capability similar heralding rates are required. The photon generation probability scales with the pumping power, but so does the conditional second-order autocorrelation, so it was not obvious in advance what an optimal choice would be. To have the option of using high circulating pump powers to push the heralding rates, a triple resonant monolithic cavity design was pursued. For robustness towards potential surface misalignment in the fabrication and due to the support of mode-size differences inherent to the large frequency difference in the pump and down-converted modes, a hemispherical cavity design is particularly suited. This design is also advantageous for the relative ease of coupling beams knowing that their foci must lie on the planar surface. Mode-matching between free space propagation and the cavity is accomplished with cemented achromatic doublets.

These were the considerations that entered into the design of the SPDC photon pair source. A remaining question is how triple resonance of the monolithic crystal cavity for pump, signal, and idler could be achieved. Generally, this requires three degrees of freedom. As the specific idler frequency is irrelevant, the pumping frequency provides one degree of freedom. Secondly, the temperature of the crystal is tuned and stabilized to yield a second degree of freedom in the crystal length. The third degree of freedom is provided by mechanical strain applied perpendicularly to the optical axis, which directly influences the refractive indices of the crystal. This final tuning mechanism is unidirectional, but with some effort triple resonance should be achievable. The conversion scheme that was settled on is pump light at approximately 404 nm down-converting into a signal at 795 nm and an idler at approximately 820 nm. All the steps of the fabrication process of a monolithic non-linear crystal cavity could be

Symbol	Property	Surface	Specified value
$L$	Length		7.0(1) mm
$\Lambda$	Poling period		10.1 $\mu\text{m}$
	Surface quality	Plane and curved	10/5 scratch/dig
		Top and bottom	20/10 scratch/dig
	Parallelism (deviation)	Plane, top, and 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) %
$R_s^p$	Reflectivity at 795 nm	Plane	91.5(5) %
$R_s^c$		Curved	$> 99.9\%$
$R_i^p$	Reflectivity at 820 nm	Plane	91.5(5) %
$R_i^c$		Curved	$> 99.9\%$

**Table 4.1:** Table summarizing the specified properties of the monolithic ppKTP cavity. The plane and curved surfaces are the end facets forming the hemispherical cavity. The top and bottom of the crystal are also polished to promote uniform straining.

outsourced to commercial enterprises. A 28 mm crystal of periodically poled KTP for non-degenerate type-II SPDC was made for us by Raicol Crystals. We had this crystal cut and polished to our specifications by Photon Laseroptik. It was then dielectrically coated by Lens-Optics. I have collected all the available specifications of this procedure in table 4.1. The source was initially set up using a 5 mm crystal, but this was short lived. Almost all the data I show and calculations I perform are therefore for our longest lived triple resonant source iteration, wherein we used a 7 mm crystal.

Let us take a quick look at how these specifications relate to our set goals. Given the cavity at 795 nm, setting the signal frequency to the Rb D<sub>1</sub> line (or near it) should be a simple matter of ensuring that this cavity is exactly resonant to the desired frequency with a spectroscopy equipped laser. More on that and setting a triple resonance point in the next section. To find the bandwidth we can take an intermediate step via the theoretical cavity finesse, which famously can be expressed as

$$\mathcal{F} = \frac{\pi}{2 \arcsin \left( \frac{1 - \sqrt{\rho}}{2 \sqrt[4]{\rho}} \right)}. \quad (4.1)$$

Here  $\rho$  is the fraction of remaining circulating power after one round-trip, i.e. in terms of the mirror reflectivities and the single-pass transmission  $T$  it is given by  $\rho = R_p R_c T^2$ . If there is no loss within the cavity, then  $T = 1$  and this expression reduces to the product of the reflectivities of the end facets. For ppKTP this is a

#### 4.1. Source Design and Expected Performance

---

good approximation in the infrared, however for the blue pumping light an absorption coefficient of  $\alpha_p = 1 \text{ dB cm}^{-1}$  was measured before the crystal was cut.<sup>2</sup> The pump transmission through a cut and coated piece of it should thus go as  $T_p = 10^{-\alpha_p L/10}$ , or  $T = 85.1(2)\%$  for  $L = 7.0(1) \text{ mm}$ . From the specified reflectivities the finesse of the signal and idler cavities are both expected to be  $\mathcal{F}_{s/i} = 70(4)$ , and for the pump, now including the measured absorptive losses,  $\mathcal{F}_p = 13.3(3)$  is expected. With this, the cavity damping rate (what I like to think of as the “single cavity bandwidth”) can be calculated with the relation  $\mathcal{F} = \Delta\nu/\delta\nu$ , so the free spectral range  $\Delta\nu$  is still required. The FSR depends on the group index,  $\Delta\nu = c/(2n_g L)$ . The group index, in turn, depends on frequency, polarization, and temperature, but can be calculated from the Sellmeier equation<sup>3</sup> [282]. At  $30^\circ\text{C}$  and with ordinary (extraordinary) polarization of the signal (idler) it is roughly  $n_g^s = 1.91$  ( $n_g^i = 1.80$ ). This yields FSRs of  $\Delta\nu_s = 11.2(2) \text{ GHz}$  and  $\Delta\nu_i = 11.9(2) \text{ GHz}$  respectively. The signal (idler) cavity damping rate is thus  $\delta\nu_s = 160(10) \text{ MHz}$  ( $\delta\nu_i = 170(10) \text{ MHz}$ ). Due to the double resonance condition the true linewidth is narrowed [389] – it is the product of two Lorentzian lines with the calculated widths. That is the way I would usually go about things numerically as well. However, a closed form for the FWHM bandwidth is provided in [389], so I reproduce it here:

$$\delta\nu_{\text{DRO}} = \sqrt{2}^{-1} \sqrt{\sqrt{\delta\nu_i^4 + 6\delta\nu_i^2\delta\nu_s^2 + \delta\nu_s^4} - \delta\nu_i^2 - \delta\nu_s^2}. \quad (4.2)$$

This final step yields  $\delta\nu_{\text{DRO}} = 110(20) \text{ MHz}$ . Further narrowing from filtering the idler is possible depending on the choice of etalon. Etalon transmissions are Lorentzian, and the same approach can be used to calculate the bandwidth after filtering. The result should then accurately capture the spectral width of the heralded photons.

The reflectivities and losses also have implications for the heralding efficiency and rates. As the heralding efficiency is defined as the probability of having a signal photon *conditioned* on the detection of an idler photon, lost idler photons do not impact it, merely their rate. The probability of a signal photon leaving the cavity through the plane surface (i.e. forwards) is estimated by [390] as

$$p_{\text{s-for}} = \frac{1 - R_s^p}{1 - R_s^p R_s^c e^{-2\alpha_s L}}, \quad (4.3)$$

---

<sup>2</sup>There is also the question of scattering losses at the cavity surfaces. Unfortunately the scratch/dig specification of surface quality provided and listed in table 4.1 is a standard that applies to the visual appearance of a surface. It specifies the maximum allowable size of scratches and pits (or bubbles) on that surface, as well as their maximum density. It is not a quantitative measure of the scattering losses to be expected, so it does not yield a way to account for them when calculating the expected performance of a cavity theoretically. Industry professionals have been contributing rants on the matter to various conference proceedings for decades [392].

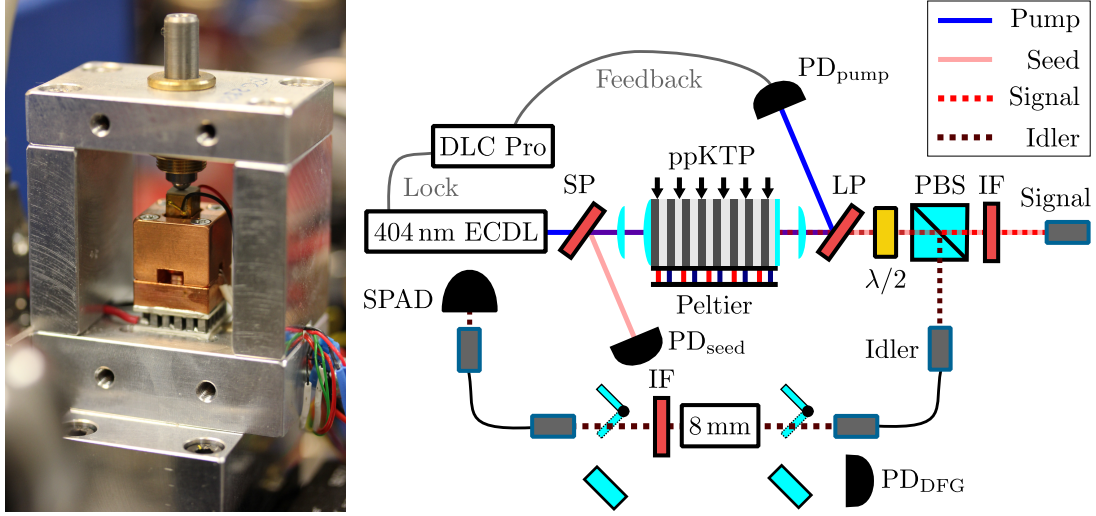
<sup>3</sup>The Sellmeier equation actually yields the refractive index  $n$ . The group index is calculated from it using  $n_g = n - \lambda \frac{\partial n}{\partial \lambda}$ .



for a signal frequency loss coefficient  $\alpha_s$ . Note that, for convenience,  $\alpha_s$  has the dimension per length in this expression as is common for small waveguides, and not decibel per length which is more common overall due to its use in fiber specification. If there are no absorptive losses, the worst case scenario for the signal loss probability (i.e. backwards emission) based on the reflectivity specification is  $p_{s\text{-loss}} = 1 - p_{s\text{-for}}(\alpha_s = 0) = 1.07\%$ . At a glance this estimate seems reasonable. The absorptive losses are too small to measure them accurately without a careful calibration of the input. We had not gone to this effort immediately, and after coating the crystal it is no longer feasible to isolate material absorption. In a similar crystal, meaning Raicol made ppKTP, others measured absorptive losses of  $< 150$  ppm/cm at 775 nm [393]. At this level absorption indeed would have negligible effect, yielding  $p_{s\text{-loss}} = 1.29\%$ . The situation is plausibly complicated by the incident pumping laser due to blue light induced infrared absorption (BLIIRA) [394]. The BLIIRA measured in [394], again with a similar crystal, is  $\alpha_s = 1.8 \times 10^{-3} \text{ cm}^{-1} - 1.8 \times 10^{-2} \text{ cm}^{-1}$  for 795 nm with intensities between  $300 \text{ W cm}^{-2} - 900 \text{ W cm}^{-2}$  of incident light at 397 nm. At this level the effects would be much more significant. For the range of absorption coefficients quoted from [394], we could expect  $p_{s\text{-loss}} = 3.6\% - 22\%$ . Again, this matter was not investigated before the crystal was coated where resulting data might have been easily interpreted. In the waist of the pumping cavity mode, which is on the order of  $20 \mu\text{m}$ , we do approach the lower end of this range of intensities under normal operating conditions. Note, however, that signs of this effect would scale with pumping power, which is not observed in experimental losses over the explored parameter space. Indeed, the limiting factor to the applied pumping power we can readily observe is the thermal stability of the crystal cavity instead.

Not all predictions based on these specifications will be perfect, in particular the finesse achieved in the infrared is only about 40. This affects the bandwidth, and indicates that the calculation for the reverse emissions may be off. Reverse emissions were not measured directly in this version of the source, but Björn Cotting measured them for a crystal with the same NIR reflectivity specifications and found that 13% of all signal photons were emitted backwards [386]. Moreover, practically unavoidable losses after the signal exits the ppKTP have not been treated at all, but will be present from the optical elements splitting signal and idler, as well as from coupling the signal into a fiber. Nevertheless, even a pessimistic outlook on the specifications and published values still seems compatible with much higher source efficiencies than are easily obtainable with external cavities or other systems, despite the lack of a reliable quantitative prediction from specifications. Finally, on to the matter of rates. The ratio between the circulating power in a cavity and the power at the input is [395]

$$I_c = T^2 \left(1 - T^2 \sqrt{R_1 R_2}\right)^{-2} I_0. \quad (4.4)$$



**Figure 4.1:** *left* A photograph of the oven housing the ppKTP cavity with the front cover removed. *right* The setup of the SPDC source including components used to establish triple resonance at the correct signal frequency. Abbreviations are SP shortpass, LP longpass, PBS polarizing beam splitter, IF interference filter, PD photo diode, SPAD single-photon avalanche diode, 8 mm etalon with given thickness.

With the measured transmission and specified reflectivities of the pumping light this yields  $I_c = 6.8(2) \times I_0$ . Combined with a diode laser input, we can be confident that this is plenty of power to crank up heralding rates to whatever the cavity temperature stability, photon quality, and any potential BLIIRA will bear.

## 4.2 Setup and Operation

The setup of the source, including elements aiding in the establishment of the desired triple resonance condition, is sketched on the right hand side of figure 4.1. A photograph of the crystal oven with the front cover removed is shown on the left. The crystal is sandwiched between two polished copper plates, with a thermoelectric cooler (TEC or Peltier) element on the bottom to control the temperature and a piezoelectric actuator on top to apply strain orthogonally to the optical axis. A screw running through the top of the mount allows for initial straining to yield the largest linear tuning range. Lenses bracket the oven to mode match collimated beams to the crystal cavity. The 404 nm pumping laser (Toptica DL pro HP) is operated with a digital laser controller (DLC pro) to enable reasonably continuous frequency tuning<sup>4</sup> (coarse adjustments require turning


<sup>4</sup>According to the manufacturer, laser diodes this deep in the blue require more sophisticated efforts to get mode-hop-free tuning ranges in the 10s of GHz than our standard analog controllers offer. A conspiracy-minded individual might interpret this as convenient.

the laser’s grating). As the crystal forms a cavity for the pump, initial resonance can be established simply by scanning the laser. Sufficient blue light leaves the plane surface of the crystal cavity to see a spectrum on the photo diode labeled  $\text{PD}_{\text{pump}}$ , which is reached by reflection at an optical longpass filter. This signal is then used to lock the laser directly to the crystal cavity resonance via feedback to the controller. To set up operation initially, the unfiltered source spectrum is recorded using a 500 mm grating spectrometer (Princeton Instruments) sensitive to single photons. This enables a coarse tuning of the crystal temperature and pumping wavelength to ensure that the strongest output cluster is on the Rb  $D_1$  line, at least to within the device’s resolution. The result is depicted<sup>5</sup> on the left side of figure 4.2. While there is a little more structure to this spectrum than the three clusters theoretically predicted this will not contribute to the heralded photon mode thanks to spectral filtering of the idler. This coarse tuning sets the crystal’s temperature in the vicinity of 33.3 °C and the pump laser’s wavelength to about 404.2 nm. Fine tuning results in small adjustments to these values. To get the SPDC process running at the precise target frequency, and simultaneously establish cavity resonance of the idler mode, a  $D_1$  line laser (TEM Lasy 795) is scanned over the cavity from the reverse direction. Initial coupling and spectroscopy is facilitated by a further photo diode  $\text{PD}_{\text{seed}}$ , this time reached via an optical shortpass. Together with the blue pump, the incident 795 nm light seeds difference-frequency generation (DFG) in the crystal, which is emitted into the same mode as the idler is in SPDC. In contrast to the idler however, this is sufficiently bright light to be detected on an amplified photo diode  $\text{PD}_{\text{DFG}}$ , reached via flip mirrors in the filtering stage, as well as with a camera (ThorLabs, DCC1645M, 5.2  $\mu\text{m}$  pixels).

The cavity spectra at the signal and idler frequencies are shown on the right in figure 4.2. With iterative adjustments to the cavity temperature, on the order 10s–100s of mK, and applied strain, yielding about 2.5 GHz total tuning range of the relative frequency, while keeping the pump laser locked, the seed and DFG peaks are overlapped at the point where the seed laser is resonant at the target signal frequency. Some narrowing in the DFG versus the width of the seed cavity mode is well visible in this plot, as predicted by equation 4.2. A Lorentzian fit to the seed mode yields  $\delta\nu_s = 370(1)$  MHz and a fit of a product of two Lorentzians to the DFG peak yields  $\delta\nu_{\text{DRO}} = 194(4)$  MHz. The latter should accurately capture the true bandwidth of the photons output by the source. Note that the value of  $\delta\nu_{\text{DRO}} = 226(1)$  MHz we reported on in [385] was not measured in this 7 mm crystal, but in the originally installed 5 mm

---

<sup>5</sup>This dataset was taken in Berlin by Janik Wolters as he had a spectrometer on hand during the initial source construction, wherein he originally installed a 5 mm crystal. No other spectrometer measurements for any sources using the triple resonant cavity design are available. In [386] Björn Cotting presents spectra of the source using a double resonant cavity design used for interfacing in chapter 5. These are notably cleaner, but I do not know if these matters are related as there are many confounding factors.



**Figure 4.2:** *left* The coarse spectrum of the source clusters, operated with a 5 mm crystal, resolved on a spectrometer. The wavelength is calibrated with a second measurement sending an attenuated  $D_1$  line laser to the spectrometer. *right* The cavity spectra of the scanned seed laser and resulting DFG at the idler frequency given by equation 2.13 before frequency filtering – recorded about 1 month after first use. The scan region captures the main mode of a single cluster. Over operational timescales of weeks the seed cavity spectrum deteriorates in an irreversible manner, the fundamental broadens and higher order mode peaks become prominent when scanning over an FSR. The DFG spectrum is pure, save for the cluster’s first side peaks at approximately  $\pm\Delta\nu$  (not shown). They have a doublet structure arising in the mismatch of the cavity’s free spectral range at the signal and idler frequencies and an amplitude about 1/20 of the main mode. At greater detunings this mismatch suppresses emissions almost completely.

crystal shortly after first set up. This is evident from the much larger FSR visible in figure 2 therein (16 GHz) than I calculated for a 7 mm crystal above. Naturally, as they are determined predominantly by material and length, and in contrast to the long list of predictions with mediocre accuracy due to known unknowns (coating, damage, etc.), the FSRs at least do behave almost exactly as expected. The numerical proximity of 194 MHz and 226 MHz is a coincidence and is not predicted by theory. Part of the reason for this is, that over long timescales there is notable variation, in particular broadening, in these modes. In day-to-day operation the effect is subtle, and it eluded us for quite a while. Unfortunately, I thus cannot present a systematic study of its variation over weeks and months of use. Instead, suffice it to say that this is one more reason the revised source design is such a boon.

A 16 h measurement of the relative peak frequencies of the DFG and seed spectra, sampling once per 5 min with a 10 MHz resolution, is used to estimate the source frequency stability. The mean absolute drift, which represents the worse case, is  $\langle |\Delta_{s,i}| \rangle \approx 15 \text{ MHz h}^{-1}$  whereas the mean itself is very close to zero. This kind of stability

is at least as good as what we observe in free-running lasers, so no frequency lock between source and memory should be required to interface them for characterization purposes.

The spatial mode of the light from DFG can be viewed with a camera. Visually, it is not an ideal match to  $\text{TEM}_{00}$  as it exits the cavity, but this is not a quantitative observation. Arguably this experiment will not resolve how the heralding efficiency later measured is limited by mode matching, induced losses, and backwards emissions, but in light of the spatial mode appearance mode-matching seems a likely candidate for the majority of the losses. A closer examination of this question will follow in the improved version of the source, so I will comment on it again in section 4.4.

Upon establishing triple resonance at the target signal frequency, the DFG signal is used to calibrate the filter stage for the idler frequency. This stage consists of an 8 mm thick etalon (274(4) MHz FWHM bandwidth at 822 nm) to remove the main emission cluster's side peaks and an interference filter (Laseroptik, 250(20) GHz (570(50) pm) FWHM bandwidth at 822 nm) to suppress emissions from other clusters.<sup>6</sup> After calibration to the DFG signal, light from a widely tunable laser (M Squared SolsTiS) borrowed for this purpose is used to characterize the filters at the idler frequency to yield the above specifications. The filter bandwidth reduces the spectral width of the mode in the same way the crystal cavity itself does. For later modeling, the final idler photon bandwidth after the filter is calculated by multiplication of the line shapes to be 150(3) MHz – its functional form is now a triple product of Lorentzians. This narrowing will present itself as a slight asymmetry in the signal-idler cross correlation. The reduction is avoidable by using a thinner etalon to remove the cluster side-peaks. The transmission in terms of idler count rate in a 400 MHz window about the peak is 68%. This too would be larger for less stringent filtering, up to about 80% maximum.

### 4.3 Output Characterization

The emissions of the source are characterized with coincidence measurements. For measuring the signal-idler cross correlation and conditioned signal autocorrelation the idler photon is detected with one single-photon avalanche diode (Excelitas SPCM-AQRH-16) and the signal photon is detected with two SPADs of the same model in Hanbury

---

<sup>6</sup>This is a little broader than the average cluster spacing of the original, 5 mm cavity measured in figure 4.2 of 410(40) pm or 195(20) GHz). As that places the flank of the IF transmission halfway to the next cluster this should make it an excellent filter, and while this measured spacing is smaller than the theory value of 270(5) GHz, following the formula of [390], this is hopefully still true for the 7 mm source as its spacing is only expected to be 30% smaller. On the failure of simple models, note that asymmetric spectra are not uncommon in real devices [288]. It is surely fair to say that we got a little lucky that model failure did not hinder us in filtering, and that we may have actually benefited from the asymmetry in requiring only a single bandpass style filter.

### 4.3. Output Characterization

---

Brown and Twiss configuration. This is the standard measurement configuration for the rest of the characterization as well, excepting the measurement of the unconditioned idler-idler autocorrelation, wherein the idler arm of the source is instead directed to the HBT. To illuminate how the measured rates relate to the heralding efficiency and pair generation rate within the cavity, which is a useful figure for modeling, I will perform the calculations with a specific example. In an experiment using triple resonance and a pump power of  $P_p = 1.2(1)$  mW measured before the cavity, signal photons are detected at a rate of  $c_s = 6.8 \times 10^4 \text{ s}^{-1}$  and idler photons are detected at a rate of  $c_i = 1.7 \times 10^4 \text{ s}^{-1}$ . The signal rate is notably higher as its mode is mostly unfiltered. The coincidence rate in an 8 ns broad window is  $r = 4.6 \times 10^3 \text{ pairs s}^{-1}$ . This rate scales linearly with the pumping power, so it is fair to give it as  $r' = 3.8 \times 10^3 \text{ pairs s}^{-1} \text{ mW}^{-1}$ . The ratio of these rates yields the efficiencies in the signal and idler arm. They are  $\eta_s = \frac{r}{c_i} = 27\%$  and  $\eta_i = \frac{r}{c_s} = 6.7\%$  respectively. The heralding efficiency is then given by  $\eta_h = \eta_s/\eta_{\text{det}} = 45(5)\%$  where  $\eta_{\text{det}} = 60(6)\%$  is the detector efficiency. It is appropriate to correct for the detection efficiency as the probability of interest is that of a signal photon being present in the signal arm upon the detection of an idler. For this, whether the signal is detected or not does not matter. The rate at which photon pairs are generated in the cavity is, naturally, the coincidence rate corrected for the losses:  $R = r/(\eta_s\eta_i)$ . Conveniently this can be expressed in terms of measured rates directly with  $R = c_i c_s/r$ , or  $R' = c_i c_s/r P_p$  when normalizing to pump power, yielding  $R' = 2.1(2) \times 10^5 \text{ pairs s}^{-1} \text{ mW}^{-1}$ . Similar heralding efficiencies are reached at all explored pumping powers, indicating that blue light induced absorption is not a significant factor in limiting this value.

The source can be operated with triple resonance for a period of a few weeks before the simple tuning procedure of adjusting strain and temperature no longer produces output at the desired signal frequency. At this point the triple resonance condition can no longer be met with the same pumping frequency. As there are visible changes to the cavity modes this is likely deterioration of the crystal due to pump-induced gray tracking [396, 397]. Multiple kinds of damage with different underlying mechanisms can occur in KTP. The most obvious kind is the production of absorbing defects,  $\text{Ti}^{3+}$  and  $\text{Fe}^{3+}$  centers.  $\text{Ti}^{3+}$  can typically be seen as gray streaks in the material – hence the name – while  $\text{Fe}^{3+}$  can produce orange dots. Another kind of invisible damage is produced by drift of  $\text{K}^+$  ions out of the irradiated area, which produces a photorefractive effect by screening the electric field [397]. This process is akin to the one that occurs in lithium niobate EOMs I mentioned in section 3.2.1. Either form of damage may be reversible by baking the crystal for a while, but I cannot report on successes in that regard. Fortunately the effect is quite localized, meaning that a new triple resonant operation point with well behaved cavity modes can be found by detuning the pump on the order of 100 GHz and readjusting the temperature and strain

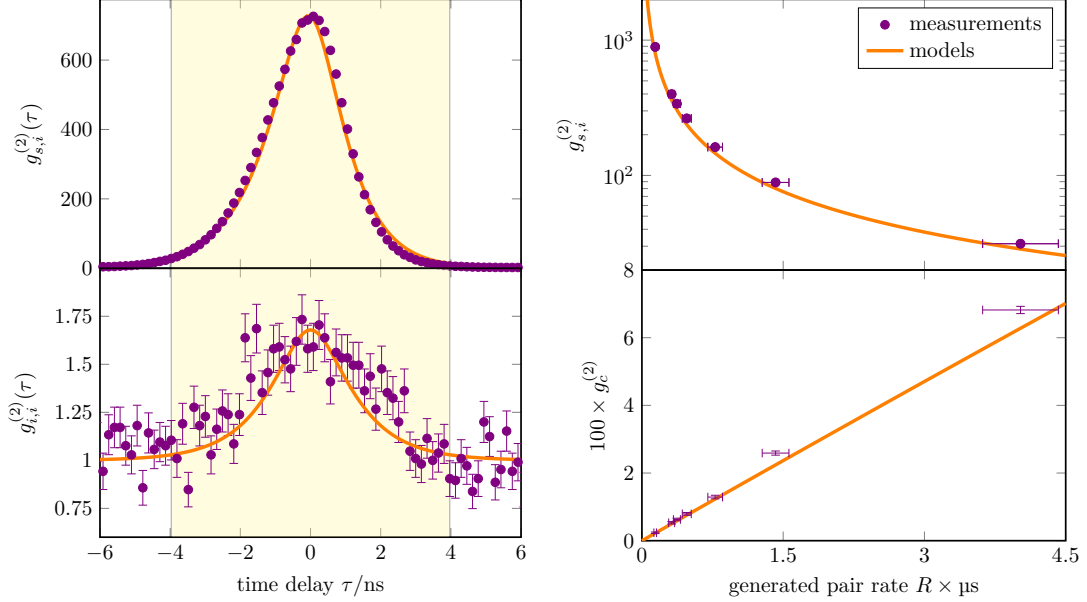
accordingly. This is quite a significant amount of work, and becomes more difficult the more it is done. After a few such readjustments the standard procedure became to operate the source without establishing perfect resonance for the pump. Sufficient pump power is available that the cavity enhancement is not necessary to get reasonable rates, so the practical condition for operation became merely that the pumping laser is lockable. To achieve a heralding rate of  $r_{h, \max} = 1.16 \times 10^5 \text{ s}^{-1}$  in this off-resonant configuration about 15 mW need to be incident on the crystal. Significantly higher rates are only stable on a time scale of 10 min. In triple resonant operation the maximal practical heralding rate is about twice as large, but this is a small payoff for the trouble it causes. For the rest of the measurements presented in this section the source is operated off-resonantly.

Second-order correlation measurements documenting the performance of the source are shown in figure 4.3. On the left, time resolved measurements of the signal-idler cross correlation and the idler-idler autocorrelation are shown. The cross correlation is measured at a heralding rate of about  $1.44 \times 10^4 \text{ s}^{-1}$  and shows pronounced super-thermal bunching for zero time delay, i.e.  $g_{s,i}^{(2)}(0) \gg 2$ . This is a first indication of the generation of non-classical light. The unconditioned autocorrelation is an excellent measure for the number of modes  $N$  present, following the relation  $g_{x,x}^{(2)}(0) = 1 + 1/N$  [242]. The total rate of heralds is kept to  $1.7 \times 10^5 \text{ s}^{-1}$  to strike a balance between low multi-photon state probability which could contaminate the mode number, and a reasonable number of coincidences. Nevertheless, the number of coincidences gathered in 142 min of integration time is only about 7000 total in an 8 ns coincidence window, which leaves fairly large uncertainties on the autocorrelation in 162 ps bins. The behavior of a single mode as a function of delay is given by [65]

$$g_{x,x}^{(2)}(\tau) = 1 + \exp\left(-2\frac{|\tau|}{\tau_0}\right), \quad (4.5)$$

where  $\tau_0$  is its coherence time. For Lorentzian light, the relationship between bandwidth and coherence time is  $\tau_0 = (\pi\Gamma)^{-1}$  [398]. At the end of the previous section the filtered idler bandwidth was calculated to be 150(3) MHz, yielding  $\tau_0 \approx 2.12(4)$  ns. As the value of 2 is reached only at zero delay, the measured value is lower even if only a single mode is present, should the detectors be subject to jitter comparable to  $\tau_0$ . The jitter of the SPADs arranged in HBT was measured and is reasonably modeled by a hyperbolic secant with a FWHM of 1.09(5) ns. For details on SPAD jitter performance and this measurement see appendix 7.2. The detector response is convolved with the expected shape of a single mode to produce the model to the measured idler-idler autocorrelation. The agreement evidences that the idler is successfully filtered to a single mode. Collecting the counts into a single 8 ns bin to yield a “full peak” estimate produces an idler-idler autocorrelation of  $g_{i,i}^{(2)} = 1.336(18)$ . In fact, this is even a little

### 4.3. Output Characterization



**Figure 4.3:** *upper left* Signal-idler cross correlation time resolved in 162 ps bins at a pair generation rate of  $0.5 \mu\text{s}^{-1}$ . The shaded region indicates the  $98 \times 81 \text{ ps} \approx 8 \text{ ns}$  region of interest in which the averaged values on the right are evaluated. This data set on the left corresponds to the fourth data point on the right. The slight asymmetry stems from the dissimilar filtering in the signal and idler arms and is predicted by the model. Note that in this plot only the line shape is modeled – the empirical peak height is simply inserted into it. *lower left* (Unconditioned) idler-idler autocorrelation time resolved in 162 ps bins. The model, described in the main text, accounts for the detector response and is compatible with a true peak value near  $g_{i,i}^{(2)}(0) = 2$ , which indicates that only one mode is heralded. *upper right* Signal-idler cross correlations evaluated for one 8 ns bin integrating the full peak as a function of the pair generation rate, effectively averaging over time. How this axis is arrived at from the detected rates is described in the text. The maximum value of the cross correlation in small bins (approximating  $g_{s,i}^{(2)}(0)$ ) at each generation rate is about 3 times larger on average. *lower right* Conditioned signal-signal autocorrelation demonstrating the high single-photon state accuracy of the source. Larger generation rates, practically equivalent to stronger pumping, come with an increased probability of multi-photon events.



larger than the model for a single mode, with or without the detector response, averaged over 8 ns, which produces the value  $g_{i,i}^{(2)} = 1.26$ . When the signal-signal autocorrelation is measured no features are visible, i.e.  $g_{s,s}^{(2)} \approx 1$  everywhere. Conversely, this confirms that filtering is truly necessary to produce a single mode. A careful accounting of spectral modes from the spectrometer measurement of the source clusters is possible in principle, and when a peak  $g_{x,x}^{(2)}(0) > 1$  can accurately be measured this approach has been shown to produce good agreement with correlation measurements in other experiments [399]. This was not the goal when the spectrometer data was collected, but a very rough estimate from the relative cluster intensities puts the effective mode number here at around  $N = 5$ . This is sure to be a lower bound as it does not correct for the substructure of the main cluster (not all of its intensity is in one mode). The detector response also enters into the model for the shape of the cross correlation. It is modeled as a convolution with an asymmetric two-sided exponential (i.e. an exponential in the absolute value of the delay, as above), rising with the  $\tau_0$  of the filtered idler, and falling with the shorter  $\tau_0 \approx 1.64$  ns corresponding to the coherence time of the unfiltered signal, calculated from the DRO bandwidth as above. The peak height is not modeled. Instead the measured peak height is multiplied with the modeled linewidth to produce the curve in the plot on the upper left of figure 4.3. The agreement confirms that the observed asymmetry is as expected from filtering, and that the bandwidth measurement of the light produced by DFG is a good proxy for the single photons produced by SPDC.<sup>7</sup>

The single-photon state accuracy is determined by measuring the conditioned signal autocorrelation. As the detector time resolution is not arbitrarily good, and the minimum time tagger bin width is not arbitrarily small, the intrinsic  $g_c^{(2)}(0)$  is technically inaccessible. This is an issue all heralded single photon sources have in common and has occasionally sparked discussions over how results are presented [400]. What is reported is necessarily a time-average of the time-dependent (Glauber) autocorrelation within a chosen coincidence window – at minimum one bin width – and the values obtained depend on this choice. This is true of the other coherence functions as well, but becomes particularly salient with the conditioned autocorrelation as a measurement of state accuracy. Baring in mind the goal of storing the photons emitted by this source, and that the sensible evaluation of memory figures of merit

---

<sup>7</sup>In chapter 3 I mentioned also using a single photon detector with better timing resolution (MPD PDM) for sufficiently accurate direct measurement of weak signals to warrant numerically Fourier transforming an arrival-time histogram to determine the bandwidth. To do so here, two such detectors would be required as both the idler and the signal need the high time resolution for an accurate transform. As no second detector was available, any measurement of single photons requires at least as many assumptions and moreover corrections as the measurement of the DFG light to extract the true photon bandwidth. I therefore hope that the accurate modeling of the photon data produced by following this route is sufficiently convincing evidence for my claims concerning the bandwidth of the photons.

#### 4.4. Limitations and Double-Resonant Version

---

takes place in a region of interest encompassing the entirety of what is read out, it behooves us not to be picky. Practically the heralding efficiency is too important to the interface; we cannot be gating photons away that are not perfectly peaked. A  $\Delta t = 8 \text{ ns}$  broad coincidence window properly capturing the full source output is therefore chosen for this evaluation. This also obviates the need for any detector response correction as it will easily be captured by the average over such a window. The trend with the pair generation rate is shown on the right in figure 4.3. Following some helpful discussions with Pavel Sekatski, we model this trend with equation 24 he derives in [401] for non-number resolving detectors. This model is expressed in terms of the photon pair generation probability. For CW pumping of the crystal this can be estimated as  $p \approx R\Delta t$  [388, 391]. Disconcertingly for a probability this estimate is unbounded, however the approximation is only valid for pumping far below threshold i.e. for small values of  $p$  to begin with. The model in [401] also includes ways to account for dark counts and detector inefficiencies. Significant dark count rates lead to larger conditioned autocorrelation values at low generation rates, see [399] for an experimental illustration of this effect. The dark count rates of the SPADs used here are low, on the order of 10 per second. This is a factor 100 lower than the smallest photon rates characterized, so this correction is negligible. Moreover the inefficiencies are already calculated out by the use of the pair generation rate. What remains is then the simple expression

$$g_c^{(2)} = 2p - p^2. \quad (4.6)$$

This model also works for the cross correlation's trend with pair generation rate. It is directly applicable exploiting a relation between the correlation functions which follows from Bayes theorem [402]

$$g_{s,i}^{(2)} = \frac{g_{i,i}^{(2)} g_{s,s}^{(2)}}{g_c^{(2)}} \xrightarrow{\text{model}} \frac{(g_{i,i}^{(2)})^2}{g_c^{(2)}}. \quad (4.7)$$

As the heralded signal entering into the cross correlation and conditioned autocorrelation inherits the idler's mode purity we substitute the unconditioned signal-signal autocorrelation with that of the idler. The conditioned signal autocorrelation remains below 1% for generation rates up to  $R_{99\%} = 5 \times 10^5 \text{ pairs s}^{-1}$ , which is an arbitrary cutoff but will allow comparison between different source iterations.

#### 4.4 Limitations and Double-Resonant Version

The characterization of the output confirms that the source's performance met the goals initially set for interfaceability. Considerable room for improvement is left open in its long term reliability and practical ease of achieving reproducible operation.

Symbol	Property	Value
$L$	Length	5.0(1) mm
$\Delta\nu$	Crystal cavity FSR	17.5(2) GHz
$\mathcal{F}$	Measured finesse (infrared)	32(2)
$\delta\nu_{\text{DRO}}$	Bandwidth	370(5) MHz
$\langle \Delta_{s,i} \rangle$	Worst case center frequency drift	$<10 \text{ MHz h}^{-1}$ typical
$\delta\nu_{\text{filter}}$	Idler filter bandwidth	2.3 GHz
$r_{h, \text{max}}$	Maximum stable herald rate	$> 2 \times 10^5 \text{ s}^{-1}$
$\Delta t$	Standard coincidence window	6.48 ns ( $80 \times 81$ ps)
$\eta_h$	Heralding efficiency	53(5) %
$R'$	Normalized pair generation rate	$6.5 \times 10^5 \text{ pairs s}^{-1} \text{ mW}^{-1}$
$R_{99\%}$	Maximum pair rate with $g_c^{(2)} < 1\%$	$7.5(8) \times 10^5 \text{ pairs s}^{-1}$
$g_{i,i}^{(2)}$	Idler-idler autocorrelation	1.28(5)
$P_{\text{pump, typ}}$	Typical pumping power	4.5 mW
$r_{h, \text{typ}}$	Typical operational herald rate	$1.5 \times 10^5 \text{ s}^{-1}$
$g_{c, \text{typ}}^{(2)}$	Typical conditioned signal autocorrelation	0.04
$\eta_{h, \text{DL}}$	$\eta_h$ after 60 m fiber	40(4) %
$\tau_{s, \text{min}}$	Minimum storage time for low noise	130 ns
$\Lambda$	Best estimate of actual poling period	10.1465 $\mu\text{m}$

**Table 4.2:** Table summarizing the performance parameters of the double-resonant cavity SPDC source. Values from [386] except for 5 penultimate lines, which are the parameters as used in and relevant to the interfacing experiments with the memory described in chapter 5 and were measured directly at the time, as well as the final entry. The estimated value of the poling period came about when I tuned the source to the Rb D<sub>2</sub> line. More details below in sections 4.5 and 4.6.


The deterioration of the triple resonance point attributed to damage required regular tedious adjustments, and within a year and a half two crystals had become completely unusable, one rather dramatically fracturing-off one of its end facets. The operational solution to the tunability problem became using the source with an off-resonant pump, which punished the triple resonant cavity design with lower rates than would have been achievable in the absence of a pump cavity. In the context of a master's thesis, Björn Cotting implemented a revised double-resonant source design using a  $L = 5$  mm piece of ppKTP without a dielectric coating at 404 nm on its curved crystal facet, but otherwise specified identically to the crystal described by table 4.1 (minimal differences are possible as the curved facet was coated in a separate run by Laseroptik but to the same specifications). In this version the pumping beam passes through the crystal twice but does not experience a resonance. This turns the pump frequency into a fully independent source tuning parameter. Double resonance for signal and idler can be established using it and the cavity temperature. The pump is locked to a

passively stable reference cavity at a freely tunable frequency by a PDH sideband offset lock. A fully detailed characterization of the source output is contained in [386]. The equivalent values to the important parameters and measurements I described for the triple resonant source are collected in table 4.2.

The figures of merit of this iteration of the source match or beat those of the previous version. Establishing the double resonance condition is considerably easier, allowing for simpler and more reliable operation. Moreover, within the year this source crystal has been in use it has not begun showing tell-tale signs of damage like a deteriorating cavity mode structure. Björn’s experiments determined that the favorable comparison of rates could be tracked back to a purer overall spectrum with fewer modes. In figure 19 in [386] he even resolves a feature in the unfiltered and unconditioned signal autocorrelation  $g_{s,s}^{(2)}(0) > 1.3$  which was not possible with the other crystal. He also measures the reverse signal emissions, finding that about 13% of the signal photons are emitted backwards, and thus implicitly laying down an absolute lower bound on the mode-matching efficiency of the source emissions into a single mode fiber of  $\eta_{m-m} > \eta_h / (1 - \eta_{\text{back}}) = 61(6)\%$  by attributing all the remaining losses to the mode-mismatch.

## 4.5 Alterations for Interfacing

Before the main results of the interfacing experiments described in chapter 5 were collected, two important alterations were made to the SPDC source to lower the rate of detected background photons in the connected systems. One background source was traced to uncorrelated photons emitted by the source at the signal frequency. These are signals that were produced in the correct mode, but were not accompanied by a heralding event. The signature of this background is thermal and has a noticeable effect on the photon number statistics of what is read out of the memory. To remove it the photon generation process is stopped at the source. As type-II downconversion is a polarization dependent process, a Pockels cell can be used to turn the polarization of the pumping laser by  $\frac{\pi}{2}$  and suppress the photon emissions to a good degree. This has the advantage of not changing the thermal load on the crystal, whose resonance condition depends on its temperature. This off-switch for the source obviously cannot be thrown before an idler photon is detected. There is therefore some time after the arrival of the heralded single photon at its destination after which this uncorrelated background persists. To find the relevant time scales in the interfaced setting, a 60 m fiber delay line is used in the source’s signal arm and a connection is made to the memory input. This delay line is motivated in section 5.2.1. Detection occurs after transmission through the memory setup, for details see figure 5.4. Having established this connection a second source of noise was found to lie in light scattered from the



**Figure 4.4:** *left* Switching behavior of the constant uncorrelated background emitted by the SPDC source when the polarization of its pumping light is turned by  $\frac{\pi}{2}$ . *right* Unless the connection between the photon source and memory includes an optical isolator, induced but switchable noise can make an efficient round trip from the memory to the source and back delaying the time full switch suppression is reached by a fixed  $\Delta t = 588$  ns.

memory setup into the fiber connecting the source. As this fiber is well aligned to a plane cavity facet on the source side, any light scattered into it will be efficiently reflected back through the memory in the signal's spatial mode. The total transmission is high enough that this background can be resolved.

Arrival-time histograms of the background from the two identified sources are presented in figure 4.4. On the left the Pockels cell based switching behavior of the uncorrelated background present in transmission through the source-memory interface is shown. The time delay is relative to the idler photon detection time. Heralded signal photons have a delay  $\Delta\tau_{i,s} = 260$  ns on this axis. The switching time is minimized as much as it conveniently can be, which means that any memory wishing to be free of this noise in the readout must have stored the photon for  $\tau_s > 130$  ns with only little room for further optimization. The programmed off-time is 500 ns. The measured suppression before the pump polarization is turned back at 900 ns is not dark count limited. Subtracting this background reveals additional structure when the memory is turned on. This is shown on the right side of figure 4.4. A switchable background feature originating in the memory operation is seen to repeat, verified by the features' joint correlation with the programmed switching time of the optical pumping beams preparing the atomic ground state for storage. Details of how this light is switched will be provided in section 5.2.3. What is interesting for the interfacing requirements of the source setup is the fact that it repeats at a fixed delay, indicating a fixed additional round trip for light of one source. The final suppression measured is limited by dark counts. The delay between the repeated features is  $\Delta t = 2 \times 294$  ns. The

#### 4.6. Tuning to the Rb D2 Line

---

fiber connecting source and memory has a measured delay of 285 ns. Beside the fiber itself,  $3\text{ m} \rightarrow \tau = 9\text{ ns}$  is a good estimate of the total free space path between the vapor cell window and the fiber on the memory setup plus the distance between the other fiber end and the planar facet of the crystal cavity on the side of the source. The vapor cell windows are wedged and angled, but fluorescence originating in the atomic ensemble during optical pumping is emitted into a large solid angle. This enables it to take this round-trip path and it is thus identified as the origin of this background. The only feasible way to eliminate the effect is with an optical isolator placed in the signal arm of the source before it is fiber coupled. Both this and the delay line itself have some impact on the heralding efficiency. In case of the isolator, the impact on the heralding efficiency results in only a small improvement in the final signal to noise ratio, so these changes are not made thoughtlessly. The total effect of both measures is about a factor 2 in SNR. The reduced value of the heralding efficiency is included in table 4.2 collecting the source specifications, as is the minimum uncorrelated noise free storage time of an interfaced memory.

To reduce the time between pair generation and idler detection, the idler filter and detection stages were compacted into a single fiber-linked module. In this configuration an external cavity was formed, delaying some of the heralding events by integer multiples of about 30 ns. As these events lead to temporal misalignment of the signal photon from its expected arrival time, a further optical isolator was placed in the idler arm. This ensures that no heralding events are systematically produced at undesired delays. The impact on heralding rate due to losses in the optical isolator is slightly overcompensated by the reduction in setup elements.

## 4.6 Tuning to the Rb D2 Line

After the interfacing results presented in the next chapter were produced, I also attempted to tune the signal photon frequency all the way to the Rb D<sub>2</sub> line at 780 nm. I found this to be possible at a pumping wavelength of about 402.9 nm, improving the best practical estimate of the poling period to  $\Lambda = 10.1465\text{ }\mu\text{m}$ , varying the sole externally specified parameter until achieving agreement with the Sellmeier equation. Note that this is merely a claim about how to get the greatest agreement between experiment and model, aiming to improve the latter's usefulness. As no uncertainty on the poling period is specified it is an obvious value to vary, but it is also conceivable that modifications to the Sellmeier coefficients could have yielded similar results. The resulting idler wavelength is 833.2 nm. Under these conditions, rotating the pump polarization was no longer sufficient to halt the pair generation process, so the Pockels cell based switch characterized in figure 4.4 was modified to include a polarizer, switching the intensity incident on the crystal. For a roughly constant

switching (i.e. heralding) rate, the rapidly varying thermal load averages well with respect to the temperature stabilization setting the cavity double-resonance condition, and does not translate to significant time dependence of the spectrum. Conversely, when the average heralding rate changes, intentionally or incidentally, the crystal set temperature needs to be calibrated for optimal output.

A circumstantially needful change in the spectral filtration of the idler photon was also made in the course of this investigation. The wavelength of the idler photon now exceeds the maximum pass-band that the interference filter used to exclude neighboring clusters for idlers at  $\lambda = 822$  nm can be angle-tuned to. This necessitated its removal, and no appropriately coated substitute was available, but I subsequently found that this had no deleterious effect on the mode number as estimated by the idler autocorrelation. This observation joins the spectrometer data taken in the course of the doubly-resonant source redesign, which is shown in figure 14 of [386], and evidences more pronounced emission into a single cluster in this iteration than what is apparent in figure 4.2.

Having mentioned these interesting modifications, and as I do not use the output of this source further in this thesis, I leave the detailed presentation of a statistical characterization of the source operated at the D<sub>2</sub> line to [383]. In brief, vis-à-vis the remaining figures of merit – efficiency, state accuracy, bandwidth, and rate – the output remains broadly comparable to the values listed in table 4.2. In particular, it is safe to say that an atomic memory operated at the D<sub>2</sub> line would be compatible with the retuned source, if it meets otherwise similar demands posed to a D<sub>1</sub> line memory under design operation.





# Chapter 5

## Zeeman Memory

*The time for compromise is over!*


Roman Schmied

---

In this chapter I will introduce a lambda memory that works towards overcoming the limitations of the implementation discussed in chapter 3 by using polarization selection rules and controlling the Zeeman states of the atoms. Many technical insufficiencies are addressed as well. Moreover, the timing logic and switching speed are completely overhauled to make the memory more generally compatible, in particular to heralded single photon sources. After discussing the overhaul to the technical implementation, I report on results obtained through interfacing the improved version of the photon pair source presented in chapter 4 with this memory. These will show that the read-out emission from the memory is dominated by single photons. After presenting the main result I detail further measurements including spectra of the memory noise, a characterization of the effectiveness of the state preparation, as well as the method used to determine the atomic temperature. In particular the source-memory interface measurements presented in section 5.3 were obtained working with Roberto Mottola. After completing this manuscript, the main results of this chapter alongside simulations implemented by Roberto were published in [382]. For details on modeling I refer the reader there or to his thesis [383].

### 5.1 Scheme

When introducing lambda schemes within the structure of the  $^{87}\text{Rb}$  D<sub>1</sub> line in section 3.1 I restricted the discussion to orthogonal linear polarizations and pointed out the problem of excited states that could absorb the signal without storing it. Implicitly, a



**Figure 5.1:** Energy diagram of the  $^{87}\text{Rb}$  D<sub>1</sub> line including Zeeman degeneracies. The signal and the control are prepared in orthogonal circular polarizations and a four-level system of the labeled memory transitions is isolated by preparing the atoms in the stretched state  $|g\rangle = |F = 2, m_F = 2\rangle$ . Only transitions to  $|e_{1,2}\rangle = |F' = 1, 2, m'_F = 1\rangle$  are allowed to the signal and  $|s\rangle = |F = 1, m_F = 0\rangle$  is set by two-photon resonance. The values of the hyperfine splittings are taken from Steck’s collection of alkali D line data [115].

second deficit of the hyperfine scheme is that interaction processes are available to the control that produce noise. Without control of the atomic Zeeman state, circular light polarizations do not seem to offer an advantage at a glance. Comparing the expected noise levels is somewhat involved due to the large number of FWM channels [403], which is only one noise source. At minimum there remain purely absorbing transitions regardless of the selected roles of the ground states. Here, these are signal transitions to  $|m'_F| = 2$ . Given such control, however, the existence of a special ground state becomes apparent. For a circularly polarized control beam with polarization  $\sigma^\pm$  the ground state  $|F = 2, m_F = \pm 2\rangle$  becomes a dark state, i.e. it is decoupled from the electric field, as no excited state with  $F' = 3$  exists on the D<sub>1</sub> line. The four-level system implied by a perfect state preparation is sketched in figure 5.1. Here the quantization axis is chosen to coincide with the propagation direction such that circularly polarized light is purely  $\sigma^\pm$ , although the sign is arbitrary assuming near zero magnetic field. This scheme has two advantages, firstly there are no non-interfering, purely absorbing pathways on which to lose the signal, and secondly the control cannot interact spuriously with the ground state atoms eliminating this noise source.

To the best of my knowledge, this scheme was first used to avoid four-wave mixing noise explicitly in an experiment combining an atomic photon source (read-only memory) and memory both based on EIT in hot ensembles [384]. In the final appendix of his

doctoral thesis, Josh Nunn credits this noise avoiding “trick” to Jean Dalibard in private communication, but goes on to dismiss its suitability in the Raman case due to fully destructive interference of the scattering pathways over the two possible excited states in the far detuned limit [198]. That the dipole matrix elements of the Raman processes with each of the excited states in this type of scheme have equal amplitudes and opposite signs holds true generally across all alkali lines, and can be traced back to the orthogonality of state components with different nuclear spin projections [404]. Indeed [404] is an interesting study as an attempt is made to quantify with what detuning this level scheme implementing the DLCZ protocol should be operated to maximize the signal-to-noise. Ultimately, a small optimal detuning range around  $\Delta = -2\pi \times 1.5$  GHz is predicted. Signal-to-noise ratio optimizations, in particular with regards to a specific process, are always system dependent: subject to temperature, the relative bandwidths of signal and noise, the scaling of other noise processes including those of the detectors, and so on. For the purposes of efficient broadband storage the increase of the necessary control intensity with detuning must also be considered. The result therefore does not transfer directly, but perhaps it can serve as a reasonable upper bound. In 2014 a study also on read-only memory was published finding that using circularly polarized beams over linearly polarized ones was sufficient to observe a reduction in FWM noise without preparation of a Zeeman state [403]. Some of the experimental choices made therein are quite strange in our context of broadband on-demand storage, for instance the interaction is tuned in between the excited states where both strongly destructive interference and considerable absorption without storage is to be expected without state preparation. The model is presented in terms of process rates per FWM channel, whereby the experimental identification of the responsible channel in each case is subject to its narrow spectrum. Once again transferring inferences to the present case without a full recreation of the work is difficult, but optimistically it seems to indicate that improvements can be expected without perfect state preparation. Following this brief review the implications of the destructive interference on the memory efficiency here remain unclear. I therefore begin with a simple model adapted to the case at hand to get a feeling for this effect.

### 5.1.1 Line Shapes and Excited State Interference

As I argued in section 3.1, destructive interference between memory interaction pathways can be addressed by choosing the detuning from the excited states so that one pathway dominates over the others. How the interaction strength varies with the detuning is described by the line shape of its transition. In hot vapor cells with buffer gas the natural  $^{87}\text{Rb}$   $D_1$  line shape, a Lorentzian peak with a FWHM of  $2\pi \times 5.75$  MHz [115], is dramatically broadened. A Gaussian contribution to the full line shape is given by

### 5.1. Scheme

---

Doppler broadening due to the atomic motion. The shift in frequency experienced by one atom due to its motion is  $\omega_D = -\vec{k} \cdot \vec{v}$  [366], where  $\vec{k}$  is the wavevector of the light and  $\vec{v}$  is the atom's velocity. Assuming a Maxwell-Boltzmann distribution of the velocities (MBD), this results in a frequency distribution with a Gaussian profile and a FWHM of

$$\delta\omega = |\vec{k}| \sqrt{\frac{8k_B T \log(2)}{m_{87\text{Rb}}}}. \quad (5.1)$$

Furthermore there is the effect of the buffer gas to consider. Via collisional processes, the presence of a buffer gas can shift and broaden atomic lines [405]. This broadening is Lorentzian in nature and its extent can be modeled phenomenologically. For the Rb  $D_1$  line  $\text{N}_2$  buffer gas induces a line shift of  $\alpha_s = -7.41 \text{ MHz Torr}^{-1}$  and a broadening of  $\alpha_b = 16.3 \text{ MHz Torr}^{-1}$  [406]. These values depend both on the species of buffer gas and on the atomic line. The shapes of these two broadening mechanisms can be convolved to yield a model of the observed line shape, a Voigt profile. The form of this absorption profile is given by [407] in units of  $(|\vec{k}|u)^{-1}$  as


$$s^I(y) = \sqrt{\pi} \Re \left( \exp \left( \frac{(a - i2y)^2}{4} \right) \operatorname{erfc} \left( \frac{a}{2} - iy \right) \right). \quad (5.2)$$

Here  $y = \Delta(|\vec{k}|u)^{-1}$ ,  $a = \Gamma(|\vec{k}|u)^{-1}$ , where  $u$  is the  $1/e$  width of the MBD

$$u = \sqrt{\frac{2k_B T}{m_{87\text{Rb}}}}, \quad (5.3)$$

and  $\operatorname{erfc}$  is the complementary error function.

To understand whether the interference affects the viability of this scheme near resonance, this line shape can be used to compare the relative transition strengths of the transitions involving different excited states as a function of detuning. We know to expect equal transition strengths for each lambda system on its resonance, as well as equal strengths in the far detuned limit. Note that the transition strengths of the lines constituting the lambda schemes are not the same: the strengths of the  $\sigma^-$  transitions ( $|F = 2, m_F = 2\rangle \rightarrow |F' = 1, m'_F = 1\rangle$  versus  $|F = 2, m_F = 2\rangle \rightarrow |F' = 2, m'_F = 1\rangle$ ) have a ratio of 3 : 1 and those of the  $\sigma^+$  transitions ( $|F = 1, m_F = 0\rangle \rightarrow |F' = 1, m'_F = 1\rangle$  versus  $|F = 1, m_F = 0\rangle \rightarrow |F' = 2, m'_F = 1\rangle$ ) have a ratio of 1 : 3 [366]. This means that working near (red-detuning from)  $F' = 1$  favors the absorption of the signal but leads to weaker control coupling and vice versa for working near (blue-detuning from)  $F' = 2$ . In figure 5.2 I first plot one side of the Voigt absorption profile of a single atomic line, normalized to the absorption on resonance. To then give an impression of the parameter regime wherein we can expect one lambda scheme to




**Figure 5.2:** *left* Absorption of a Voigt shaped atomic line normalized to peak absorption as a function of detuning  $\Delta$ . *right* The ratio of the transition strengths (i.e. line shapes) of two Voigt profiles separated by the excited state splitting. The plot is anti-symmetric about the center frequency between the excited states. A value of 1 on the y-axis is a valid proxy of fully destructive interference between the lambda schemes and is approached asymptotically on either side of the plotted range.

dominate over the other, I plot the ratio of two such Voigt line shapes detuned from one another by the excited state splitting. The parameters entering this figure are an atomic temperature of  $T_{\text{at}} = 50^\circ\text{C}$  and an  $\text{N}_2$  buffer gas pressure of  $p_{\text{N}_2} = 5$  Torr. Note that y-axis of the right plot does not map neatly to memory efficiency – it is not quantitatively representative of the two-photon lambda process – but it does indicate that concerns about interference significantly hindering this memory scheme’s effectiveness are unnecessary within moderate detunings to the red (blue) of  $F' = 1$  ( $F' = 2$ ).

## 5.2 Setup

Required components to practically implement this memory scheme break down into the same categories I identified in section 3.1: a (proxy) signal, the control, an atomic ensemble, means of state preparation, and noise filters. The new target signal is, of course, the topic of chapter 4, and the optical setup of the laser test signal remains as I described it in section 3.3.1. A photo of the memory itself is shown in figure 5.3. This corresponds to the bottom part in the sketch of the complete and combined source/memory setup depicted in figure 5.4. Since the experiments of chapter 3, considerably more space became available for these experiments. In the course of this the setup has spread out, and little effort has gone into compacting its total footprint. I will detail modifications and improvements to the components here, focusing on

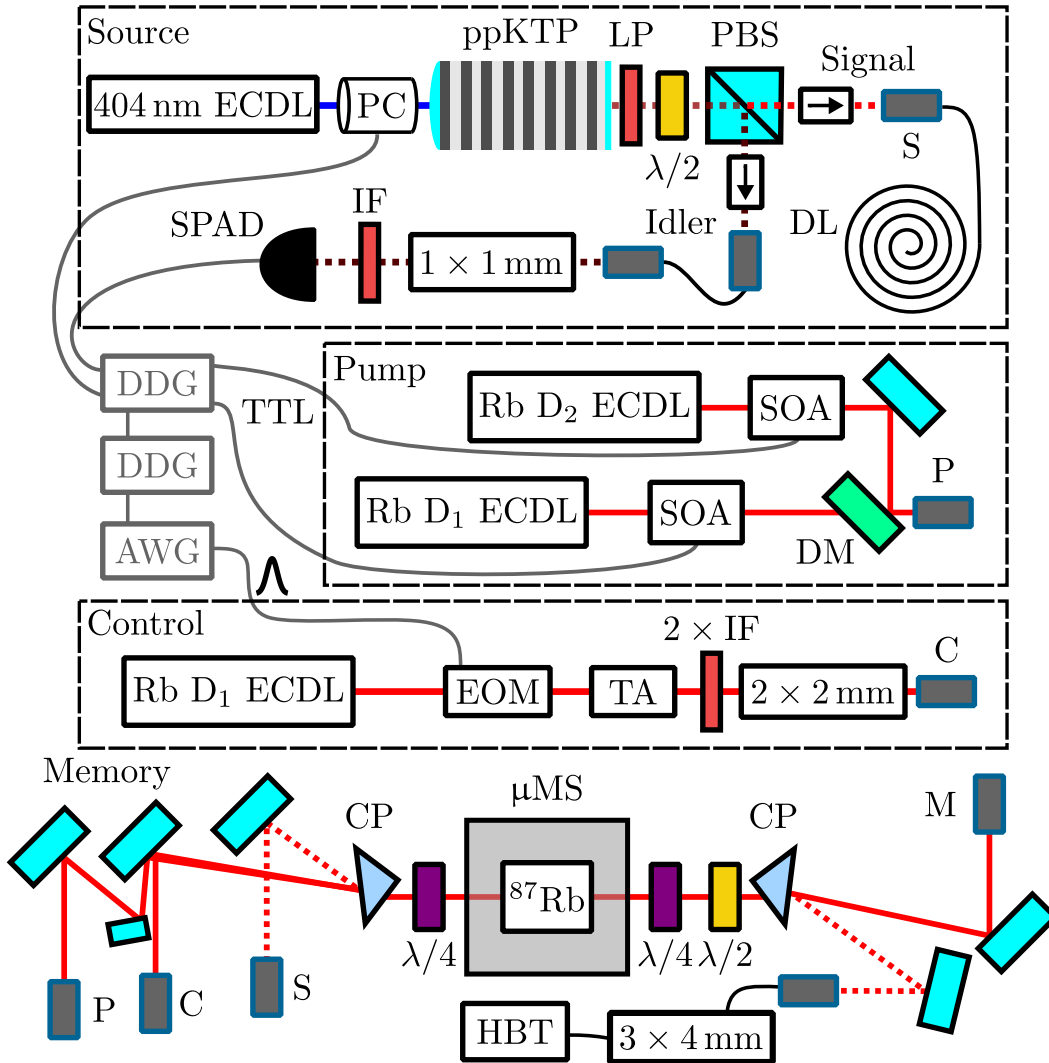


**Figure 5.3:** A photo of the heart of the memory experiment, which is now spread out over a whole optical table.

the differences to what I described previously, and addressing the technical wish list laid out in section 3.6. Some components, such as the electro-optic modulators for generating pulses from CW light, are identical and thus I do not describe them again. The most significant technical change overall, demanded by the goal of storing photons from a heralded source, is that the entire memory process must now occur on demand upon a spontaneous external trigger. This constitutes a paradigm shift with regards to electronic requirements, in particular, insertion delay, jitter, and the maximum response rate to external triggers become significant figures of merit. Optical switches also need to be reactive. Further, state preparation of a Zeeman state requires one laser more than state preparation of a hyperfine state, and it will also turn out that switching requirements preclude the use of AOMs. I will now quantify these timing requirements and describe the electronics that meet them, then discuss the optical setup.

### 5.2.1 Requirements of an Externally Triggered Memory

Experiments with a deterministic single-photon source, deterministic as always in the sense of section 2.2.1, can be repeated periodically. A convenient consequence is that if a device takes a while to respond, it can just be triggered earlier. With periodic repetition this is even possible when the delay is longer than the entire experiment time. Experiments storing high bandwidth signals for short times approach this limit quite easily, indeed it is reached in the hyperfine memory. The situation is markedly



**Figure 5.4:** Experimental setup: ECDL external cavity diode laser, PC Pockels cell, ppKTP monolithic periodically poled potassium titanyl phosphate cavity, LP optical longpass, PBS polarizing beam splitter, IF interference filter,  $x \times y$  mm  $x$  etalons of thickness  $y$  – see table 3.1, SPAD single photon avalanche diode, DL 60 m fiber delay line, DDG digital delay generator, AWG arbitrary waveform generator, SOA semiconductor optical amplifier (fiber connections omitted), DM dichroic mirror, EOM electro-optic modulator (fiber connections omitted), TA tapered amplifier, CP calcite polarizing prism,  $\lambda/2$  half-wave plate,  $\lambda/4$  quarter-wave plate,  $\mu$ MS 4-layer mu-metal magnetic shield, HBT Hanbury Brown and Twiss configured single photon detectors, M fiber connection for monitoring the control. The labels S, P, and C represent the fiber connections of the signal, pump, and control to the memory respectively.

different when attempting to interface a memory with a heralded photon source. Upon the detection of the idler photon the signal photon is already on its way. Famously, light is pretty fast, and because the start of the experiment has only just been marked there is no “earlier” when an irresponsive device could have been triggered. The time between the entry and exit of a signal into a device or medium is called *insertion delay*.<sup>1</sup> The insertion delay of free space in this sense is about  $3.3 \text{ ns m}^{-1}$ , and that of optical fiber is about  $5 \text{ ns m}^{-1}$ . The first order of business in interfacing these systems is therefore to ensure that the insertion delay of the signal photon to the memory is longer than that of the idler photon to its detector. Fancier means of doing so are mentioned in section 2.1.1, but optical fiber is simple to use, minimally variable, and reasonably priced. But how much should the signal be delayed by?

Let us first reexamine the hyperfine memory with this newly gained perspective on responsiveness. The electronic switching was performed by a general purpose AWG (Tektronix AWG7122C, described further in section 3.2.2). I measured the insertion delay of this device to be 162 ns, which I intuit to be reasonable performance for such a complicated device – compare the  $10 \text{ ns m}^{-1}$  delay caused by mere coaxial cable – but is quite long in absolute terms considering the experiment timescales. The optical switching of the control/pumping beam included an acousto-optic modulator. Here the exact values vary with the device, its driver, and the cable length, but a typical figure between electronic trigger and the beginning of the optical response is 400 ns. Disconcertingly, a closer investigation of the matter reveals that the time from the beginning of the optical response to when a suppression exceeding  $1 : 10^3$  is achieved,<sup>2</sup> even when using a focused beam for improved switching speed, is much longer – on the order of a few microseconds. The effect of this may be small for storage times of 50 ns, but to ensure that the memory is not read out preemptively by a randomly scattered photon during longer storage a delay of this order is required. Unfortunately, the impact of this approach on the interface performance would be deeply negative.

For single photons it becomes particularly apparent that optical fiber suffers from  $4 \text{ dB km}^{-1}$ – $5 \text{ dB km}^{-1}$  transmission losses due to absorption at 795 nm. A 100 m fiber, for example, corresponding to a delay of just 500 ns, maximally transmits 90(1) % of the signal, even with perfect incoupling. By the time an AOM finishes switching at a level of  $1 : 10^3$  a fiber delay line would have absorbed  $> 50 \%$  of the signal, and thus truly entered the realm of Pyrrhic solutions. An alternative was found, but not by way of a device available off-the-shelf at the time. I will detail it in section 5.2.3.

---

<sup>1</sup>This term is also used when the signal entering a device is a trigger and the device output is arbitrary, including the case of an electronically triggered optical response, although *response time* might have been a more fitting name in these cases.

<sup>2</sup>This ratio does a much better job of capturing the time taken to reach the desirable attenuation of a strong beam, say with  $I \approx 10I_{\text{sat}}$ , into the weak probe regime,  $I < I_{\text{sat}}/100$ , than commonly specified 90 : 10 fall times.




Yet another matter is the memory lifetime. Interference experiments between stored photons and photons directly emitted by the source, or between subsequently stored photons, are of principle interest as a means to determine the memory fidelity. For reasonable coincidence rates in such experiments, as well as for a fair characterization of the memory performance, the delay must be shorter than the lifetime. The results of the hyperfine memory experiment have plenty of room for improvement in that regard, but set a baseline that encourages caution. Finally, longer delays also increase the minimum total experiment time, which sets the maximum repetition rate. If this rate does not exceed the heralding rate of the photon source the memory becomes the limiting factor to the interface's speed. All these factors motivate my initial choice of a 60 m fiber between the photon source and the memory, which turns out to delay the signal by 285 ns.<sup>3</sup> Anti-reflection coatings on the end-facets enable a transmission of about 93 %, including coupling.

As I originally discussed in section 2.1.4, the stopping of light with EIT relies on the control field being turned off while the signal is within the storage medium. Concretely this means that the optical signal and control pulse, or technically its falling flank, must coincide in the vapor cell temporally. The limit to the quality of this temporal overlap is given by the shot-to-shot variance in the timing signal delays, which is called *jitter*. The most common specification of jitter is one standard deviation (RMS) of a series of delay times.<sup>4</sup> The FWHM depends on the distribution of the jitter, for the normal distribution it is larger by a factor  $2\sqrt{2\log 2}$ . The maximum time difference, or peak-to-peak jitter, represents the worst case scenario – usually this is roughly 5 times the RMS value. In chapter 3 the electronic signal and control pulses were both generated by the AWG7122C, and I begin this evaluation implementing storage as therein. The total internal jitter of this AWG is specified to be about 30 ps RMS, which is thus the fundamental limit to the quality of the time-alignment. Practically, as this is much smaller than the pulse widths, the alignment can be fully optimized. In the scenario of storing a heralded single photon, the trigger to the memory is produced by a single-photon detector measuring an idler photon. The single-photon avalanche diode used for this purpose (Excelitas, SPCM-AQRH-16) has its time resolution specified to  $> 350$  ps. The measure used for this specification is ambiguous. While purely electronic devices are almost exclusively specified with RMS jitter, a full width at half maximum measure would be typical for specifying optical impulse response. Moreover

---

<sup>3</sup>To determine the absolute delay of the fiber I measure the relative delay in signal detection with and without the signal passing through it.

<sup>4</sup>Jitter as a measure is intended to capture timing variance due to phase noise. The time interval over which it is measured should therefore be short. Longer term variation, driven by temperature, is called *wander*, and can include sources like thermal length variation in coaxial cable. A third detrimental effect, called *walk*, is variance in the amplitude of the trigger signal, which translates to timing variance in when the device's trigger voltage is reached.



**Figure 5.5:** In a pulse storage experiment the programmed delay between the electronic signal and control pulse generation is scanned and the retrieved pulse area is recorded. Lower values of  $\tau_{\text{gen}}$  correspond to the control pulse arriving earlier than the signal, but do not map perfectly onto the arrival time of the optical signals at the vapor cell due to some optical path length difference between the amplification and attenuation stages. The errors of the data correlate as they are caused by an offset.

slight misalignment or insufficient focusing of the incident light rapidly degrades the detector’s temporal response [408]. To verify its performance empirically an optical signal significantly shorter than 350 ps is required. A characterization of this kind can be found in appendix 7.2 and is most compatible with the specification being an RMS value, but is confounded by potentially sub-optimal focusing. For the following estimates, I will assume it to be an RMS value. This is dubious, but ultimately irrelevant – all models for data acquired with these SPADs use empirical values for the achieved time resolution, and by time of writing we had acquired detectors with better performance and unambiguous specifications we could verify independently. Either way, when storing heralded photons detector jitter fundamentally limits how well the signal and control can be time-aligned, and jitter relative to this trigger from “downstream” devices involved in the control pulse generation adds further variance to their coincidence in time.

To quantify the effect of temporal misalignment of the control to the signal, the delay can be varied in an experiment not limited by jitter. A scan of the relative generation time by the AWG of the signal and control pulses used in chapter 3 is shown in figure 5.5. Signal pulses, attenuated to the weak probe regime (peak power of a few nanowatts) for high SNR and a rapidly performed, thus well controlled, experiment, are stored, and the retrieved pulses are measured on a sensitive avalanche diode (MenloSystems APD210). The variation in the retrieved pulse area is modeled as Gaussian with a FWHM of 3.6(3) ns. This is the scale on which jitter drastically reduces the memory efficiency. The data are slightly asymmetric, and could instead

be modeled with independent sensitivities to early and late control pulses. This has a plausible physical explanation in the asymmetry of the signal pulse, which is still programmed to produce the shape shown in figure 3.4. The right flank of the data, which corresponds to the control pulse arriving later than optimally, leaves the control reasonably well overlapped with the tail end of the signal. It falls about 30% slower than the left flank, where the faster rising edge of the signal lies. Obviously the sensitivity also depends on the control pulse width, shorter pulses requiring better alignment. An experiential rule of thumb I have found in experiments not limited by jitter is that alignment in steps  $1/20$  of the control pulse width are sufficient to reliably find a small deviation from the maximum memory efficiency with regards to time-alignment. For visualization, this is 1.5 times the narrower step-size used about the peak of the data shown in figure 5.5. This sets the bound on jitter having no significant detrimental effect on memory efficiency, and it appears to be reached in an experiment interfacing a heralded source just by the idler photon detector performing ideally. It is therefore desirable to produce as little further jitter in the memory as possible.

The intrinsic jitter in a device's output is not necessarily the same as its jitter to an external trigger. Electronic synchronization can accomplish wonders for timing accuracy, but a device reacting to a probabilistic process must, by definition, do so asynchronously. Optimization seems simple. Measure the trigger signal on an oscilloscope, find the steepest point on the voltage flanks, set the device's trigger level to the voltage measured at this flank (assuming the device allows for this), and enjoy minimal additional jitter. This does not generally guarantee that the jitter is now the same as it would be if the device were acting as its own clock, which is sometimes the only figure directly specified. While absent from the documentation included with the AWG7122C, I found a technical brief made available online wherein the manufacturer mentions asynchronous operation and estimates the jitter to an external trigger to be 500 ps, dependent on the sampling rate [409]. The measurement thereof is not detailed further, and it is ambiguous whether this is intended to be a peak-to-peak or RMS value. Fortunately, the additional jitter induced by a device to the signal chain can be measured by synchronizing an oscilloscope with the original trigger and measuring the temporal variance of the device output. This measurement can of course also be performed with all the devices processing a signal in series to yield the system jitter. If  $k$  devices in a signal chain have uncorrelated jitters  $\delta\tau_i$ , the system jitter is given by the root sum square  $\delta\tau_{\text{tot}} = \sqrt{\sum_{i=1}^k \delta\tau_i^2}$ . When measuring jitters it is important to ensure that the oscilloscope is actually capable of resolving them, and to note that jitter usually increases with delay. In this way, and using a high sampling rate ( $40 \text{ GS s}^{-1}$ ), low jitter ( $< 2 \text{ ps RMS}$ ) oscilloscope (LeCroy WaveRunner 640Zi), I measure that the AWG7122C has an RMS jitter of about 130 ps, concluding that the specification

probably represents a peak-to-peak value. Starting with 350 ps detector jitter, this would only increase the expected RMS system jitter to 373 ps.

Finally the question arises whether the maximum rate of external triggers a device accepts could be a limiting factor. Usually this rate is the inverse of the duration of the output plus a small constant. Bizarrely, the AWG7122C used previously as well as another device of the same model found in-house are both subject to a known manufacturing defect, which limits the accepted rate of external triggers to about  $f_{\max} = 9 \times 10^3 \text{ s}^{-1}$ . I got this confirmed by Tektronix support after failing to exceed this rate in a series of ever simpler tests, but alleviating it would have required returning the device for repair. This slightly exceeded what I wanted to do with a device that I had merely borrowed, and with the knowledge I have detailed in this section in mind I searched for suitable specialized devices to replace it with. I will briefly cover these devices' novelties, but their main purpose remains just as described in section 3.2.

### Logical Signals and Pulse Generation

While the true “master clock” of an experiment interfacing a heralded single-photon source is the idler photon’s detector, all devices in the memory will require a timing signal, i.e. a voltage pulse, to react to the presence of the signal photon. The device directly triggered by the SPAD should therefore serve to deterministically delay and redistribute trigger or follower signals as they are required. Here I opt for the Highland Technology T564, a 4-channel digital delay generator (DDG in figure 5.4). In particular with regards to the timing figures discussed in this section this device performs well. Its insertion delay is only 20 ns, making it compatible with a short signal delay line. Its jitter to an asynchronous external trigger is specified to  $< 35 \text{ ps RMS}$ , well below the jitter produced by the idler’s detector. The device manual contains a much more detailed specification as a function of delay [410], as well as a bunch of notes on what limits the device in what delay range. For a programmed output duration  $t$ , its maximum rate is  $f_{\max} = (t + 60 \text{ ns})^{-1}$ , which is an insignificant increase in the minimum experiment time. Its resolution is 10 ps. In pulse train mode it can also repeat the output of one or all of its channels periodically after the programmed sequence ends for the first time, with a delay variable in steps of 20 ns, minimum 60 ns. Its 4 channels are not quite sufficient to address every device, and two T564s are required to operate the Zeeman memory. An associated benefit is that its compact size enables strategic placement near where the signals need to go to cut down on cable length and the associated delay. Using these two devices in series I also perform a jitter measurement here, finding 41(2) ps RMS jitter for a delay of 200 ns at 200 kHz. The error represents short term drifts in the standard deviation. Assuming equal contributions from the two devices, each contributes about 29 ps RMS, which agrees perfectly with the specification

in [410].

The positive experiences with Gaussian control pulses qualify the demands for AWG timing resolution, as detailed pulse shaping seems not to be needed. For pulse generation I therefore turn to the PicoQuant PPG512 (AWG in figure 5.4). This device integrates a digital-to-analog converter and amplifier to produce up to 12 V with 8 bit amplitude resolution and 200 ps time bins. The insertion delay is small, on the order of 10 ns. The jitter is loosely specified as  $< 200$  ps. Measuring, I find 95(2) ps RMS jitter on the electrical signal (92 ps subtracting the trigger jitter). The maximum rate is set by the inverse of the device record length, which as its name suggests is 512 samples or 102.4 ns. This produces a slight oddity of operation. If the storage time exceeds the record length, the device must be programmed to produce one pulse and receive two triggers. This is facilitated by the delay generator's pulse train mode. While it is supposed to be possible to have a trigger interrupt the PPG512 and restart its output, my device does not support this feature (it was implemented in the first hardware revision). Therefore, to produce two control pulses within the record length for short storage times they must be programmed into one pulse sequence. For longer storage times, the step size of the T564s pulse train mode restricts variation to 20 ns increments. To achieve arbitrary storage times a second delay channel would have to be used, but this is not necessary for characterization. In a real use scenario the photon would presumably be read out on-demand, with a trigger coming from the external device requesting the photon. Provided with such a trigger, the memory operates fully asynchronously.

The final signal chain for the control pulse consists of the idler detector, two T564s in series, and the PPG512. In a further measurement using one T564 for trigger synthesis as above but also including the second T564 and the PPG512 I find the optical signal<sup>5</sup> after amplification jitters by 126(3) ps RMS. This is about 20% more than can be accounted for from the electrical signal, and corresponds to an additional uncorrelated source with 75 ps RMS jitter. Due to the small effect on the final value I have not investigated this matter further. Subtracting the DDS jitter from the optical measurement and replacing it by the specified 350 ps photon detector jitter yields a predicted RMS jitter of the optical control pulse in an interfaced experiment of 371 ps, subject almost entirely to achieving the detector specification. To estimate the effect of this amount of jitter on the obtainable efficiency I model it as a Gaussian line shape (unity area) with  $\sigma = 371$  ps and convolve it with the line shape of the

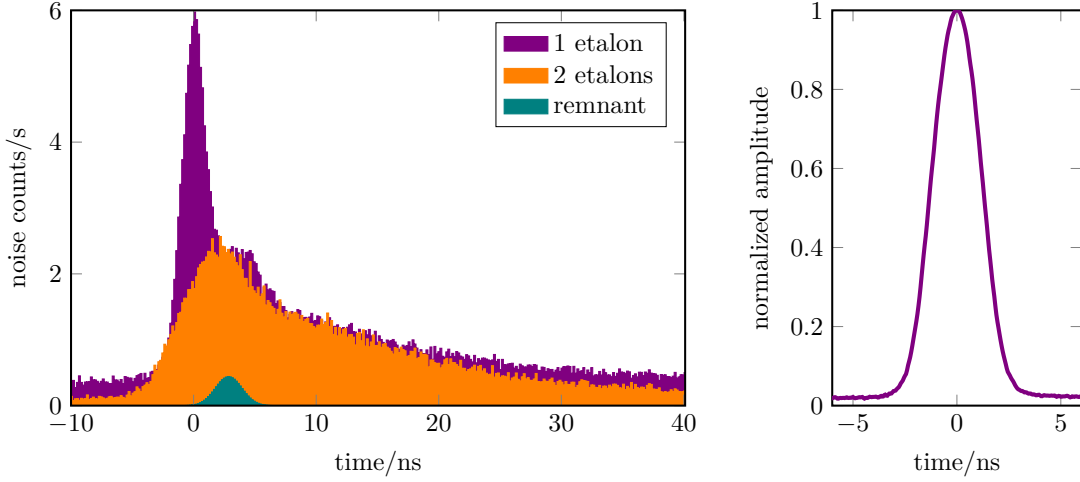
---

<sup>5</sup>The optical detection is performed with a fast InGaAs photo diode (ThorLabs DXM30AF, DC–30 GHz) with a specified impulse response of 19 ps. This measure is the FWHM of the minimum detector response, i.e. the response to arbitrarily short signals. If the shape of the response is also specified the full instrument response function is known, and its standard deviation is equivalent to a jitter specification. While the shape is not explicitly stated we can judge from the small FWHM of the response that it should have no significant impact on the measured jitter value.

model of figure 5.5. A subsequent comparison of amplitudes predicts a 3% relative reduction in memory efficiency due to jitter. Note that this prediction depends both on a signal/control shape dependent model and on achieving a best case detector specification, which is not directly verified. As the photon waveform and control pulse in the final series of storage experiments are both somewhat narrower than what enters the (quasi) jitter-free model, this is a lower bound on the detrimental effect of jitter on our interface.

### 5.2.2 Control Pulse Amplification

The generation and amplification process of the control pulse is depicted in the box labeled control in figure 5.4. An electro-optic amplitude modulator (Jenoptik AM785) is used to generate pulses from a CW diode laser. Then these pulses are then amplified with a tapered amplifier and spectrally filtered before being sent to the vapor cell. The pulse width can be varied using the PPG512; the standard test case is programmed as a 4 ns FWHM Gaussian pulse. While the principle approach is the same as described in section 3.3.2, there are a number of notable technical improvements to remark upon. Firstly, the control laser is now exclusively used to produce control pulses. This means that the EOM can be locked by the low duty cycle technique described in section 3.2.1, which is convenient and mildly beneficial to its performance. The EOM is the same model as that of chapter 3, but the device has been replaced once within 5 years due to cumulative photorefractive damage. Consequences of this damage included ever faster drifts of the required DC bias at full optical power, on the order of 20 s to cover the full range. The maximal transmission also suffered, decreasing by around 30%. To improve behavior and spare the replacement EOM a similar fate the laser power is attenuated to 5.7 mW before the EOM, leaving a peak power of 800  $\mu$ W at the TA input. The drift in the required modulator bias is significantly slowed by this change, and the device remains within one locking range for hours. This lower power is sufficient because the TA gain chip was also replaced (Toptica BoosTA pro with TA-790-2000-1 chip). The gain region of this TA is peaked at 789 nm with a 10 dB-below-peak bandwidth of 38 nm, making it easily capable of its full performance at 795 nm. This also obviates the need for anamorphic beam shaping optics before the amplifier, in fact a circular beam seems to work better with this chip than an elliptical one. The TA is not saturated by this light; the CW output is around 1.1 W of a maximal 3 W. This is done intentionally after experimentally determining the required control intensity, so that excess power is reduced before the EOM and amplification. In the rest of this characterization data the generated control pulses correspond to those that optimize memory performance. Direct comparison of memory experiments conducted with the same control intensity on the atoms but with attenuation either



**Figure 5.6:** *left* The necessity of spectral filtration of the TA output for effective suppression after the memory is illustrated by the beneficial effect of additional filtration on the noise rates. The transmission of the control pulses by the second etalon is about 75 %, which is divided out of the 2 etalon rates. The falling flank of the second trace is well fit by an exponential function with a decay time around  $\tau = 17$  ns, in good agreement with what was observed using spatial filtering. The peak looks like it might still contain a Gaussian remnant – this component is isolated by subtracting the exponential fit to model potential control leakage. *right* The shape of the control pulse programmed to be a 4 ns FWHM Gaussian is measured using a fast diode (ThorLabs DXM30AF, DC–30 GHz). The plot shows a measured, normalized pulse almost perfectly described by a Gaussian with a FWHM of 3.77(4) ns. This small discrepancy between programmed and measured width is reproduced for other settings and found to be constant, in absolute terms, over the explored range (programmed FWHM < 8 ns).


after the TA or before the EOM reveal noticeably better performance in the latter case, both with regards to memory efficiency and noise. Potentially this could be due to both devices. Less stress is put on the EOM lock leading to a lower background, and the TA operates in the linear regime yielding more accurate pulse shapes. I have not quantified how much this is the case and which of the devices benefits more from the arrangement, but I have observed increased noise from the TA when saturated in a separate comparison, and the improvements in the EOM stability are evident.

Increased power output from the TA also mandates more thorough spectral filtering before the memory. To estimate the suppression of the signal frequency in a control pulse I take a  $-40$  dB expected ASE noise floor from the TA and subtract the suppression at the hyperfine splitting  $\Delta_{\text{hf}} = 6.8$  GHz, as listed in table 3.1, of each etalon used to prepare the pulse. This value can be compared to the expected suppression of control frequency light in the spectral filters after the memory. Leakage of a control pulse

through these filters is easily identified in arrival-time histograms. On the left side of figure 5.6 I show a comparison of the control pulse induced noise in a periodically repeated experiment with either one 2 mm etalon (expecting  $-58$  dB suppression) or two (expecting  $-76$  dB suppression) cleaning its spectrum before the memory. An etalon of this thickness induces about 1 ns of additional delay, which is compensated electronically. In both cases two interference filters (Laseroptik, FWHM  $\approx 400$  pm at 795 nm, angle tunable) are also used to filter the light, but transmit the control and signal frequencies equally. The light is sent to the vapor cell, but unlike previous experiments no misalignment between the spatial modes of signal and control is intentionally produced. Any reduction in transmission of the control over the signal in this stage can be traced back to a minimal difference in the quality of the mode matching to the fiber following the memory due to slightly different beam waists. After the memory, however, three 4 mm etalons (expecting  $-78$  dB suppression) serve to remove light not at the signal frequency. In the one etalon pre-filtering case the noise has an obviously Gaussian peak component. This is removed by adding the second etalon, leaving a fairly pure exponential decay. A small, roughly Gaussian remnant appears to remain near the peak, but this is ambiguous as any sharp onset of light could also appear as a Gaussian peak due to the instrument response of the detector. Subtracting an exponential model based on the long tail of the noise reveals thorough misalignment in time of the remnant from a peak that would be caused by leaking control light, which is centered at zero in all cases. Its amplitude, compared directly, is also incompatible with the measured CW etalon suppression at a detuning equal to the hyperfine splitting by nearly an order of magnitude. These observations support an atomic origin of this light, but the effective removal of the coherent control light will ultimately be verified spectrally.

Due to the low duty cycle operation, the average incident power on the optics is significantly lower when the TA produces pulses than what it would be if it were seeded in CW. Initial alignment is done at low power of equivalent average, but the final optimization of the etalons and the fiber coupling to the memory are all done directly with pulses to ensure optimal performance in the operational conditions. For the control pulses used for storage the etalon transmissions are about 75% and the first fiber coupling efficiency peaks around 65%. To determine the Rabi frequency the peak control power on the atoms needs to be determined. As a CW characterization is not a reliable proxy for pulsed performance at low duty cycles, using it as a stand-in for the peak pulse power as in section 3.3.2 would be subject to a significant systematic error. Therefore I generate a pair of control pulses, as in a read-in/read-out attempt, deterministically at  $f_{\text{rep}} = 200$  kHz (comparable to a heralding rate, save for the periodicity) and measure the integrated power  $P_{\text{int}}$  with a thermal power sensor (ThorLabs S401C). The idea is to find the pulse energy by  $P_{\text{int}} = 2E_{\text{pulse}}f_{\text{rep}}$ , and use





**Figure 5.7:** Due to the finite extinction ratio of the EOM, the intensity in the control mode is non-zero when the control pulses are “off”. To quantify this amount of light the strongly attenuated output of the control fiber to the memory is measured on a single-photon detector while the trigger logic of WCP storage for 160 ns at 200 kHz is applied. The suppression of spurious control light settles just under  $-28$  dB below peak at long time scales, but shows considerable dynamics right after the control pulses are generated. The origin of these dynamics are ringing in the EOM after the strong voltage pulse.

the measured temporal pulse shape shown on the right side of figure 5.6 to extrapolate the peak power from the area. This measurement has some uncertainty associated with it due to fluctuations. The ambient background can be subtracted, and variations can be minimized by darkening the room. Nevertheless, the measured integrated control power cannot be determined better than  $P_{\text{int}} = 720(40) \mu\text{W}$  due to random variations over time. This is not necessarily a reflection of variations in the peak amplitude, as there is also the background from the finite extinction ratio of the EOM to consider.

As the duty cycle is only about 0.45%<sup>6</sup> the energy contribution of the background would be roughly equal with a pulse suppression of  $-26$  dB. Based on the EOM specification the scale of the effect is large enough to warrant measurement. Further, while operated in the linear regime, it must also be verified that the TA does not significantly alter the extinction. Due to the need for simultaneous time resolution and contrast ratio strong attenuation (about ND100) is placed after the spectral TA filtration and the fiber is redirected to the single photon detectors. The result is shown in figure 5.7. While the long term suppression of  $-28$  dB would suggest only  $110 \mu\text{W}$  of power originates outside of the pulses, there are considerable dynamics near the control


<sup>6</sup>I’m estimating this figure here by treating 3 times the FWHM of the control pulse as the on time. This time window contains  $> 99.9\%$  of the pulse energy.

pulses. Directly comparing the counts in two  $3 \times \text{FWHM}$  wide regions around the pulse peaks with the total counts reveals that only 42% of the total light actually lies in the control pulses. The best available estimate for the pulse energy is thus 760(40) pJ. The area of a Gaussian peak with amplitude  $P_{\text{max}}$  is given by  $E_{\text{pulse}} = \sqrt{2\pi} P_{\text{max}} \sigma$ . From the measured pulse shape we have  $\sigma = 1.60(4)$  ns yielding a peak power of 190(20) mW. Subject to realignment, the maximum possible peak control power on the atoms, achieved when saturating the TA, is about 700 mW. At a CW output around 2 W, achieved with about 15 mW of light incident on the EOM, the bias drift remains relatively slow and the peak power on the atoms reaches around 450 mW.

### 5.2.3 Switching of the Optical Pumping Mode

This section discusses the components of the box labeled pump in figure 5.4. For the memory to be ready to store a photon upon the detection of an idler photon, it is necessary to prepare the ground state between storage attempts. As I will show in section 5.4, the state preparation is short lived. To reject as few herald triggers as possible it is therefore desirable to switch off the state preparation in response to the idler's detection, rather than opening an acceptance window for heralds after a fixed time of state preparation. To do this effectively, an optical switch satisfying the timing requirements discussed in section 5.2.1 is needed. As mentioned therein, the AOM previously used for this task cannot be transferred to this scheme due to intolerably long insertion delay. An alternative is to pass CW laser light through a semiconductor optical amplifier (SOA) and rapidly switch the current through its gain medium. At 1550 nm such devices are well known and already highly advanced in terms of extinction ratios and gate integration [411]. The principle of operation of a SOA is as in a tapered amplifier, but without the taper. Current is passed through a semiconductor chip with a broad spectrum, whereupon a small amount of seeding laser light incident on the chip triggers a chain of stimulated emission at its frequency. When no current is passed through the chip, not only is there no gain, but the medium also absorbs seed light. Without the charge carrier densities associated with tapering the gains and currents are lower, and the complete device can be fit in small, fiber-pigtailed butterfly packages without temperature control becoming a problem.

Fast switching of small currents is (allegedly) quite straight forward, and as SOAs are threshold devices the demands on the extinction of the current are not extreme. A device capable of switching up to 200 mA in a few nanoseconds was kindly designed and built for me by our in-house electronics lab (internally documented as SP 1047). It is switched by follower TTL logic, provided directly by the delay generator, and the on time is continuously variable up to CW operation with the voltage pulse duration. It also integrates temperature stabilization for the SOA by attaching a TEC-controller



**Figure 5.8:** *left* The suppression achieved by using a SOA as an optical switch on experimentally relevant time scales. *right* Photograph of the SOA switch, mounted on a base plate also holding the PCB driver and temperature controller.

(Meerstetter Engineering, TEC-1091). For consistent operation the SOA is held at 25 °C. I use this current source in combination with a fiber-integrated SOA capable of outputting 20 mW (Superlum SOA-332-DBUT-PM). This limit is imposed by the damage threshold of the fiber connections. The optical seed to the amplifier can be provided by a diode laser, so gain is not required to achieve full output, and the parameter can be freely varied to optimize switching performance. A photograph of the device is shown on the right hand side of figure 5.8. In the interim since these devices were assembled as switches, some manufacturers have begun selling operationally similar integrated devices as optical switches at Rb wavelengths. I currently have no experience with such commercial devices, but the design has a long established history in that role at telecom wavelengths. This could be a sign of the latest transfers of commercial high speed means of optical control from their telecom origins into the frequency range of alkalis, akin to what happened with waveguide EOMs.

On the left hand side of figure 5.8 I show a characterization measurement of the SOA suppression as a function of the delay to the falling flank of the follower signal, while the SOA is continuously seeded by 650  $\mu$ W of light from an ECDL. The current is switched between zero and 132 mA at 1 MHz with a 50 % duty cycle. Hoping to maximize absorption in the off-state, I operate the SOA with about 15 dB gain to produce 20 mW at the output. The output is sent to a single photon detector<sup>7</sup> and attenuated until the detection rate is similar to the repetition rate to exclude saturation or latching. In

<sup>7</sup>For this characterization I used a superconducting nanowire single photon detector (Single Quantum EOS), acquired after the rest of this chapter's data were collected. On our SPADs, the high extinction was washed out by a long exponential tail I can only attribute to detector behavior for signals with

## 5.2. Setup

---

comparative measurements I observe the switching behavior and achieved suppression vary slightly with the set current if it is near the SOA threshold. The transmission without current is no longer accurately measurable with a power meter, implying an ultimate extinction of at least  $-75$  dB. The rising flank of the light closely follows the 10 ns current rise time with minimal overshoot. The following on time is uneventful and not shown. Upon switching, after a short insertion delay mostly due to coaxial cable, the intensity falls rapidly, again closely following the 3 ns current fall time. The final suppression shown in figure 5.8 is limited by dark counts. For an experiment repeated at  $f_{\text{rep}}$  integrated for  $t_{\text{int}}$ , with a bin width  $t_{\text{bin}}$  and a detector dark count rate  $f_{\text{dc}}$ , the expected number of dark counts per bin is  $N = f_{\text{rep}} t_{\text{int}} f_{\text{dc}} t_{\text{bin}}$ . Integrating for 6 min and using  $48 \times 81$  ps  $\approx 4$  ns bins so that none are empty, the detector dark count specification of  $f_{\text{dc}} < 10 \text{ s}^{-1}$  yields an expected number of dark counts per bin of  $N < 11.7$ . In the data shown in figure 5.8, the mean count number per bin during the off time is only 8.3, leading to the preceding conclusion. This performance was reproduced in all four of these devices we assembled, both at the Rb D<sub>1</sub> and D<sub>2</sub> wavelengths, and I choose to show the data with the smallest dark count rate. In some datasets, the resolved suppression seemed minimally worse, but a complete lack of dynamics lets me believe that this is appropriately attributed to dark count rates minimally exceeding the specification. Dark counts have many possible sources, including black-body radiation and detector illumination itself [412], and their rates can plausibly vary a little from day to day without obvious changes having occurred. Remaining light coming out of the SOA during the off time consists of transmitted seed light and is therefore spectrally pure. If even greater suppression proved necessary, it could be achieved using a saturable absorber, like a hot vapor cell.

Optical pumping to the Zeeman state  $|F = 2, m_F = 2\rangle$  requires the light of two frequencies: *pumping* depletes  $|F = 2, m_F \neq 2\rangle$  and *repumping* depletes  $|F = 1\rangle$ . While there are a few options for the repumper, the pumping beam must be  $\sigma^+$  polarized setting the polarization of the mode. As the  $\sigma^+$  transition  $|F = 2, m_F = 2\rangle \rightarrow |F', m'_F\rangle$  is not dark on the Rb D<sub>2</sub> line, we should stick to D<sub>1</sub> for this laser. In principle, it would now be possible to modulate an  $\omega = \Delta_{\text{hf}}$  sideband onto a single pumping laser, for instance with an electro-optic phase modulator, to address both ground states simultaneously. In practice, if the modulation is phase coherent, it is likely that this will produce undesirable dark-states via coherent-population trapping [194]. For this reason, and for fully independent switching, the repumping is performed with a second laser. For efficient combination and separation of spatial modes a laser on the Rb D<sub>2</sub> line, addressing the transitions  $|F = 1, m_F\rangle \rightarrow |F' = 2, m'_F = m_F + 1\rangle$ , is used. After both lasers pass through SOAs to enable their switching, the beams are combined on a

---

such a high dynamic range. As it is not used elsewhere in this thesis I point the reader to [383] for more information on the detection system in particular and to [412] for a general review of SNSPD.

dichroic mirror (Semrock RazorEdge LPD02-785RU-25 with specified transmissions  $T_{D_1} = 97.4\%$ ,  $T_{D_2} = 0.52\%$ ) and the combined mode is fiber coupled. The power ratio of the pumping and repumping beams on the atoms is about 2 : 1, optimized on memory noise performance, with a pumping power of about 18 mW after the fiber.

### 5.2.4 Vapor Cell and Memory


The vapor cell at the heart of the Zeeman memory is a Precision Glassblowing made, 75 mm long quartz cylinder with 19 mm outer diameter and wedged windows. The wedge of the windows is  $2^\circ$  and they are attached to the cell at  $11^\circ$ . It contains enriched  $^{87}\text{Rb}$  ( $< 1\%$   $^{85}\text{Rb}$  specified), 5 Torr of  $\text{N}_2$  buffer gas, and paraffin anti-relaxation coating on the walls. By equation 2.8, this nitrogen pressure extends the laser atom interaction time by about a factor 4. The cell is housed in a heater enclosed by a 4-layer magnetic shield (Twinleaf MS-1L, specified shielding factor  $10^6$ ). The heater (Sacher Lasertechnik) can maintain atomic temperatures up to about  $70^\circ\text{C}$ . For the experiments in section 5.3 the atomic temperature is  $50(1)^\circ\text{C}$  and the optical depth on the resonance of the signal transition is about 25. The determination of the temperature via spectroscopy is described in section 5.4. The signal and control modes, exiting their respective fibers in orthogonal linear polarizations, are combined on a polarizing calcite prism. To guarantee no reduction in the polarization extinction ratio due to strain-induced birefringence, this is not a glued prism pair as is commonly used to compensate achromaticity. Instead it is a single, right-triangular prism. This prism was sold to me by the optical workshop Bernhard Halle upon request. It is one half of what they usually sell as a glued Wollaston polarizer.

This is a somewhat unusual shape for an optic used in transmission, and as single-photon signals are at hand it is worth briefly checking what the Fresnel equations have to say about how much light is lost to reflection in the absence of an anti-reflection coating. The Fresnel coefficients of reflection for  $s$  and  $p$  polarizations and generic incidence and refraction angles  $\theta_i$  and  $\theta_r$  respectively are,

$$r_s(\theta_i, \theta_r) = -\frac{\sin(\theta_i - \theta_r)}{\sin(\theta_i + \theta_r)}, \quad (5.4)$$

$$r_p(\theta_i, \theta_r) = \frac{\tan(\theta_i - \theta_r)}{\tan(\theta_i + \theta_r)}. \quad (5.5)$$

Transmission through the prism is subject to two interfaces. For the air-calcite interface let  $\alpha$  be the angle of incidence and  $\beta$  be the angle of refraction, related by Snell's law as  $n_1 \sin \alpha = n_2 \sin \beta$ . Here  $n_1 \approx 1$  is the refractive index of air and  $n_2$  is the refractive index of calcite. As calcite is birefringent, it has a different refractive index for ordinary (s-polarized) and extraordinary (p-polarized) rays. They are  $n_o \approx 1.65$



**Figure 5.9:** *left* Visual aid to analyzing the transmission  $P_2/P_1$  through a right-triangular prism with the Fresnel equations. *right* The expected maximum transmission of each light polarization as a function of the angle of incidence on the prism.

and  $n_e \approx 1.48$  for light at 795 nm [413]. Then, for the calcite-air interface, the incident angle is fixed at  $\gamma = \frac{\pi}{4} - \beta$  by the prism geometry, and the angle of refraction  $\delta$  relates by  $n_2 \sin \gamma = n_1 \sin \delta$ . See the left side of figure 5.9 for a visualization. The maximal (intensity) transmissions through the entire prism for s and p polarizations are then  $T_{s,p} := P_2/P_1 = (1 - |r_{s,p}(\alpha, \beta)|^2) \times (1 - |r_{s,p}(\gamma, \delta)|^2)$ . These values are plotted on the right side of figure 5.9. The maximal transmission for s-polarization (vertical in the laboratory frame) is just under 80% for an angle of incidence  $\alpha \approx \frac{2\pi}{9}$  rad. For a weak signal that is a lot of loss. For a strong control it is a potentially hazardous amount of stray reflection. Either way, this is an optical element that should definitely be anti-reflection coated. As the goal is to combine the signal and control modes it is not appropriate to attempt optimizing the incident angles separately. Instead, for initial alignment at least, a diagonally polarized beam should be sent through the prism in reverse, whereupon the incidence angle can be optimized for the more sensitive polarization by rotating the prism.

Upon combination of the modes on the polarizer, a quarter-wave plate turns the orthogonal linear polarizations into orthogonal circular polarizations. After the memory a further quarter-wave plate reverses this process and a half-wave plate is used to compensate for any constant rotations. Then a second calcite prism is used to separate the control from the signal and each mode is coupled into a fiber. The control mode is monitored with a fiber-coupled photo diode (ThorLabs DET025AFC), while the signal mode is sent to the spectral filtration stage. The combining prism induces a small amount of astigmatism, but correct orientation of the splitting prism after the cell compensates this effect quite well. The lens coupling the signal is carefully chosen and

aligned to mode-match the diverging beam to yield around 83% transmission through the fiber.

Due to the large vapor cell size, and the magnetic shield limiting optical access to it, optical pumping must be implemented on-axis. The polarization requirements therefore mean pumping ideally shares a spatial mode with the control laser. To avoid the losses associated with combining the beams, the atoms are pumped under a small angle. Further, as in the past experiment, the signal and control modes are focused into the vapor cell directly by the lenses in their fiber outcouplers. To see what is going on at the position of the vapor cell, I set up a beam profiler (DataRay, WinCamD-UCD12, 4.65  $\mu\text{m}$  pixels) accessible via flip mirror in an equivalent distance and on a rail to cover the length of the cell. This way the signal (control) is focused to a  $e^{-2}$  diameter of 480(6)  $\mu\text{m}$  (520(6)  $\mu\text{m}$ ) in the center of the vapor cell. These waists are sufficiently large that the size of the beams varies by only a few percent over the length of the cell. The peak resonant Rabi frequency of the control can now also be calculated with equation 2.4, where the dipole moment of the control transition ( $|F = 1, m_F = 0\rangle \rightarrow |F' = 1, m'_F = 1\rangle$ ) is  $\mu_{s,e} = \frac{1}{\sqrt{12}} \times 2.54 \times 10^{-29} \text{ C m}$  [115]. This yields  $\Omega = 2\pi \times 400(30) \text{ MHz}$ . The pumping mode is collimated, with a  $e^{-2}$  diameter of 1 mm. A D-shaped mirror is used to overlap it with the control under an angle of 2.95(15) mrad (extrapolated geometrically). As the chromatic dispersion of the calcite prism is not compensated, the spatial mode of the repumping beam is horizontally displaced by about 40  $\mu\text{m}$  with respect to that of the pumping beam in the center of the vapor cell. While noticeable at the scale of the 15 nm difference in wavelength of the D lines, this chromatic displacement is negligible for frequency differences on the order of  $\Delta_{\text{hf}}$ , and thus does not impact the alignment of the signal and control modes.

### 5.2.5 Filtration

In chapter 3 we saw the detrimental effect of spatial filtration on memory lifetime in hot atoms. Now, it is not the ambition of this experiment to produce a particularly long lived memory. I presented other authors' methods that have proven to be successful at extending the memory lifetime in EIT memories to millisecond and even second scales at the end of section 2.1.4, and would argue that with sufficient technical effort they are transferable to this memory as well. However, if the memory performance relies on spatial filtering, and the memory lifetime is thus limited by rapid spin wave dephasing, then this argument is clearly wrong. This situation motivates implementing noise filtration that relies only on polarization and spectral filtration, without intentionally inducing any misalignment between the signal and the control. By  $E = h\nu$ , the extrapolated control pulse energy of 760(40) pJ implies that an average control pulse contains about  $3 \times 10^9$  photons. This appears to be a manageable amount, but there

are some caveats. Firstly, even in CW, measuring how well light is extinguished at the involved scales is difficult. No single detector provides the dynamic range that would be required for a direct measurement, least of all with a linear response. Precision would therefore require the careful calibration of a multitude of detectors and attenuators to cover the range from powers around 1 W, all the way to single photons. A much simpler route which I follow in this section is to characterize the suppression of the control polarization, or frequency respectively, in CW and element by element. This approach is flawed. The CW suppression of the control frequency does not necessarily capture the performance for a broadband pulse, and the measured CW suppressions of individual elements do not generally add up neatly to yield the total suppression of a filtering system.

Nevertheless, the actual goal of a filtration system is to remove a lot of noise with as few elements as is possible to keep signal transmission high. Ideally it leaves only noise behind that cannot be filtered as it shares both the polarization and the frequency of the signal. Whether this is accomplished is a question with a binary answer, and does not rely on measuring the exact suppression directly. There is also the question of what exactly remaining noise consists of, as filtration may also help with excess photons produced by an atomic response. The matter of optimizing and quantifying filtration is therefore more complicated fundamentally than just determining the suppression of the control laser. These measurements benefit from the context of a concrete storage experiment. I therefore defer a more informative spectral analysis of the noise to section 5.4, and just cover the case of CW control suppression here.

The intrinsic polarization extinction ratio of the calcite prisms used for combining and separating the bulk of the control and signal is expected to be incredibly high and fully symmetric. As they consist solely of single pieces of calcite, the fundamental limit lies in the material purity itself. Beyond the concerns of a detector's dynamic range, the leaked power must be measured in a sufficient distance to ensure the measurement is not limited by randomly scattered light. To ensure that I capture their performance as it is in experimental conditions, I measure the control power after the fiber bringing the signal to the spectral filters. Rotating a half-wave plate between the prisms I can modulate the CW control power measured with a photo diode power sensor (ThorLabs S130C) after the fiber between about 80 mW and  $< 500$  pW, limited by the detector range, yielding an extinction ratio better than  $10^8 : 1$ . This measurement also includes the two quarter-wave plates but leaves the atoms cold. As mentioned previously, I spent some time comparing ways to mount polarizing prisms for optimal performance. I tested small amounts of epoxy, soft adhesive from a glue stick, and nail polish, as well as clamping and taping. The winning method is using a tiny piece of double sided tape, ideally small enough to raise serious concerns about whether the prism is attached to the mount at all. To notice the difference, the prism's intrinsic polarization extinction



ratio must be very high, but this is the only method<sup>8</sup> by which I have ever observed extinction ratios better than about  $10^{6.5} : 1$ .

The spectral filtration stage consists of three 4 mm etalons in series. What I have said about etalons in section 3.3.4, as well as table 3.1 summarizing their individual performances, applies here as well. The combined FWHM bandwidth of all 3 etalons measures around 300(10) MHz and the transmitted spectrum is well modeled by the cube of the spectrum of a single etalon. On their alignment, it is common practice to align cavities of all kinds by scanning a laser over them. When used in series, as these filtering etalons are, this is not feasible for the etalons after the first one. It is important to begin the alignment of the now constant frequency laser to the subsequent etalon by centering its spatial mode on the etalon and overlapping the back-reflection. Iterating this procedure over a sufficient distance eliminates alignment errors save for those between the surfaces of the etalon itself. The etalon temperature and transmission are not reliable measures until the fundamental mode is addressed in this manner, as it is very easy to “get trapped” in the local transmission maximum of a higher-order mode when relying on it to guide the alignment. Once this process is complete the etalon temperature can be scanned and set to maximally transmit the target frequency. Only then should the transmission be optimized on directly to find the best external alignment compromise given the inherent tilt. After completing this process for every etalon in series, scanning a laser passing through all of them reveals the combined bandwidth and spectral shape. It is possible that some external cavities, i.e. free-space cavities between the etalons, have become sufficiently well aligned to impart a modulation to the spectrum transmitted through the filters. In this case I have found that small tweaks to the alignments can remove the modulation without significantly altering the total transmission. In this way I obtained peak transmissions of 90 % through 2 of the three etalons and 80 % through the third. After the third etalon another fiber coupling brings the signal to the detectors. The transmission of a few milliwatt CW signal from before the fiber leading to the etalons to after the fiber leading to the detectors reaches about 48 %. The maximum total transmission of such a signal through the entire setup that I have measured is 36 %, although 30 % is a more typical and stable value. As before this serves as a lower bound on the losses single photons are subject to.

---


<sup>8</sup>The only method in which I was the one mounting the optic, that is. I have also used ready-mounted Glan-Thompson polarizers from Bernard Halle selected for high extinction ratios which have performed to their  $10^8 : 1$  specifications. Unfortunately, these are effectively linear polarizers with no way of combining two orthogonal polarizations. It is not possible to improve polarization filtration using such devices if the initial combination/separation of signal and control is performed on a worse polarizing beam splitter. On the polarization error scale at hand, these effectively act as depolarizing beam splitters!

### 5.3 Single Photon Storage


In this section I will present the results obtained and figures of merit determined by interfacing the Zeeman memory and the SPDC pair source. I will begin by condensing the previous section into the experimental sequence. The time between the detection of an idler photon and the arrival of the signal photon in the vapor cell is about 270 ns, as it takes about 15 ns for the idler to be detected after the signal enters its delay line. In a storage and retrieval experiment the memory is first initialized by the pumping and repumping lasers, setting the minimum experiment time of 2  $\mu\text{s}$  during which the detection of idler photons is rejected on a hardware level. This means the ground state gets prepared for at least 1.5  $\mu\text{s}$  between storage attempts. The value chosen here is determined by comparing the results obtained in a separate data set without hardware rejection beyond the 500 ns storage sequence, i.e. no enforced pumping time at all, wherein the effect of longer or shorter minimum pumping times are compared by post-selecting. A significant detriment in the evaluated figures of merit is only observed when reducing the rejection time to  $< 1 \mu\text{s}$ , but 2  $\mu\text{s}$  is chosen to ensure similar conditions for all experiments. Further, this time remains short enough to have only minimal impact on the heralding rate.

After this minimum duty cycle of the state preparation, the detection of an idler photon triggers a DDG (Highland Technology T564) to switch off both the atomic pumping and repumping beams as well as the pumping of the source as detailed in section 4.5 until after the photon is retrieved. The repumping beam is left on 50 ns longer to minimize the occupation of  $F = 1$ . Optical switching is prioritized to minimize noise, as it takes some time for the excited state to relax, and the cables are kept as short as possible. This reactive configuration can be cast as a memory with some dead-time after readout, which in return remains ready to accept a photon any time afterwards. The trigger is also relayed to a second DDG for less delay critical tasks, including triggering the generation of control pulses and time stamping the idler photon detection with a time-to-digital converter (qtools quTAU). The typical rate of accepted heralds in these experiments is  $1.5 \times 10^5 \text{ s}^{-1}$ , which is almost identical to the rate of generated heralds. The heralding efficiency after the 60 m fiber delay line is  $\eta_h = 40(4) \%$ , and the conditional second-order autocorrelation of the signal photons at this rate is measured to be  $g_c^{(2)} = 4.21(2) \times 10^{-2}$ .

The memory is operated at a detuning from  $F' = 1$  of  $\Delta = -2\pi \times 700 \text{ MHz}$ . An arrival-time histogram for repeated attempts of storing a heralded photon for 160 ns is shown in figure 5.10. The detection is preformed with two SPADs (Excelitas, SPCM-AQRH-16) in HBT configuration. Data are collected for 20 min of integration time. The other two curves represent different estimates of the noise performance. The source blocked scenario represents the typical noise estimation performed for memories,



**Figure 5.10:** Arrival-time histograms for storage experiments and noise characterizations showing the regions of initial photon leakage through the memory and the retrieval after 160 ns. The data are histogrammed in 162 ps bins. As no intermediate features are visible on a linear scale the time-axis is broken to show only these peaks. The shaded area marks the 6.48 ns ( $80 \times 81$  ps) wide region of interest for the retrieval for which all figures of merit in the text are specified.



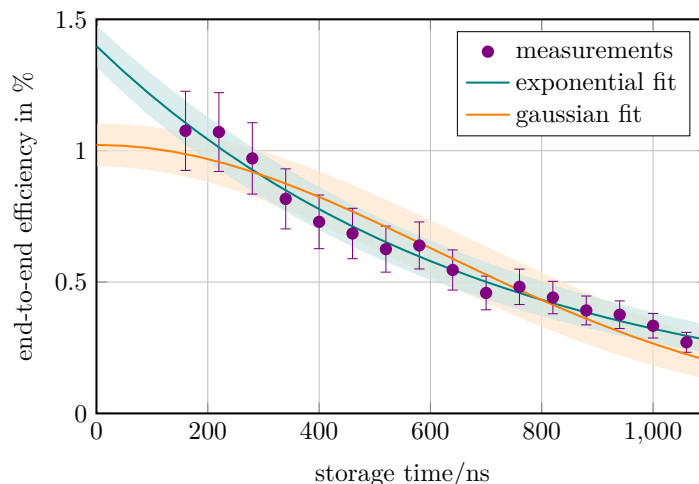
**Figure 5.11:** Arrival-time histograms for storage experiments and noise characterizations on a logarithmic scale. The data are histogrammed in 1 ns ( $12 \times 81$  ps) bins. The  $y$ -axis is in counts and somewhat arbitrarily normalized to the peak of the no read-in curve in the retrieval window (8327 counts). Several features too small to see on a linear scale, labeled A–F, can be identified and are discussed in the main text.

### 5.3. Single Photon Storage

---

including what I showed in chapter 3. It is the amount of noise produced when the input is physically blocked. This serves as a comparative measure to other memories as it characterizes the memory in isolation, but it does not capture the total amount of noise in the interfaced system accurately due to noise originating from the source – see also section 4.5. An alternative measure is shown in the no-read curve. Here the setup is left just like it is in the storage trials, but the first control pulse writing the photon into the atoms is omitted. The same data are plotted logarithmically to reveal the features in the comparatively small noise in figure 5.11. In the negative delay region the idler photon has already triggered the electronics, and we see exponentially decaying fluorescence from the atoms after the pumping beams are switched off. Here A marks where the pumper switches and B marks where the repumper switches. In the region around C, after the photon is stored, the discrepancy between the storage and no-read curves corresponds to unintentional read-out of the stored photon. It is caused by ringing in the EOM switching the control and its limited extinction ratio. The feature present in both curves right beneath the label at around 30 ns is due to afterpulsing of the SPAD. This is a kind of false detection event distinct from a dark count in that it occurs systematically in SPADs one detector dead time after the detection of a photon, with just under 1% probability in our detectors. A similar feature can be seen in the storage curve after the read out peak. Potential causes of afterpulsing are diverse and the effect is difficult to model generally [414]. The steep feature labeled D corresponds to the Pockels cell switching off the photon source. As this is done as quickly as possible this motivates the choice of a 160 ns storage time. At earlier times an additional background of uncorrelated photons from the source is present. The remaining discrepancy between the no-read and source blocked curves estimates the secondary effect of the first control pulse on the atoms. The rising feature labeled E corresponds to the pumping lasers being switched back on after the retrieval is complete. The source switching back on is labeled F.

When the memory is run, in a read-out window of 6.48 ns (shaded region), and for a total of  $N_{\text{herald}} = 1.60 \times 10^8$  storage attempts,  $N_{\text{ret}} = 4.54 \times 10^5$  photons are retrieved. When the source is blocked the number of memory induced noise counts for the same number of experiments is  $N_{\text{noise, mem}} = 2.91 \times 10^4$ . This noise is induced by the control and is quite well fit by an exponential as in figure 5.6, however as its peak rate has been reduced to  $< 1 \text{ s}^{-1}$  (in 162 ps bins) by state preparation the measured decay time is even shorter. This yields the signal-to-noise ratio of the memory in isolation as  $\text{SNR}_{\text{mem}} = 14.6(20)$ . To capture the total noise including the source I perform a measurement omitting the read-in control pulse. This induces a small systematic error as the atomic response to the read-out pulse is also influenced by the read-in pulse – note the slight difference between the first and second control pulse induced peaks of the source-blocked curve. Alternatively, it is possible to apply both



**Figure 5.12:** Memory efficiency evaluated as above for various storage times to measure the lifetime. These measurements stem from the day after I collected data suitable for statistical evaluation, and were all obtained in one day to minimize random variations. The integration time for each measurement is 5 min. The efficiency was not optimized as thoroughly here and turned out to be about 15% lower at 160 ns storage time. I have scaled the data by this factor for an accurate internal efficiency calculation of the main data. The exponential fit yields a  $1/e$  lifetime of 680(50) ns and an initial efficiency of 1.40(8)%. The Gaussian fits considerably worse. It estimates the lifetime at 860(80) ns and the initial efficiency at 1.0(1)%. The shaded regions represent the 95% confidence intervals of the fits.

read-in and read-out pulses, but to misalign the read-in from the signal in time. This then necessarily measures the total noise at a slightly different (hypothetical) storage time. Comparative measurements reveal only small differences in the results of these approaches within the read-out window. The noise counts measured in this “no read-in” curve’s region of interest are  $N_{\text{noise, tot}} = 3.86 \times 10^4$  yielding a total signal-to-noise ratio of  $\text{SNR}_{\text{tot}} = 10.8(15)$ . With these noise counts I determine the end-to-end efficiency to be  $\eta_{e2e} = 1.1(2)\%$ , dividing out  $\eta_{\text{det}} = 60(6)\%$  and  $\eta_h = 40(4)\%$ . As previously the quantum efficiency of the single photon detectors, specified as 60(6)%, is the dominant source of uncertainty. The end-to-end (internal) efficiency of the unintentional read-out from leaking control light during the storage time is 0.38(5)% (1.26(17)%). The signal-to-noise ratio bodes well for the quality of the retrieved photons. Indeed, the conditioned autocorrelation of the retrieved photons is  $g_{c, \text{ret}}^{(2)} = 0.177(23)$ , confirming that the memory emission is dominated by single photons. The statistical error is given as  $1/\sqrt{N}$ , where  $N = 57$  is the number of triple coincidences in the region of interest. This number is too small to subdivide the region of interest into smaller bins and give a more time resolved characterization of the  $g^{(2)}$  with reasonable uncertainties.

### 5.3. Single Photon Storage

---

A measurement of the memory lifetime is shown in figure 5.12. The exponential  $1/e$  lifetime of the memory as measured by the drop in its efficiency is 680(50) ns. This is shorter than the atom-laser interaction time  $\tau \approx 6.4 \mu\text{s}$  equation 2.8 predicts due to the use of a buffer gas. It is better, but not well, predicted by the ballistic estimate of the transit time  $\tau \approx 1.5 \mu\text{s}$ . The 5 Torr  $\text{N}_2$  we use is quite low compared to examples in [166] explicating the time scales of different motional regimes, however the expected Rb- $\text{N}_2$  mean free path of  $13 \mu\text{m}$  is considerably smaller than the beam waists, leaving little justification for treating the atomic motion as ballistic. In contrast to the hyperfine memory, the lifetime data are fit considerably better by an exponential decay than by a Gaussian flank. This is compatible with processes causing a loss of atoms limiting the lifetime, i.e. the atoms move out of the spatial mode of the control before read-out, get their spin scrambled by a collision, or are prematurely read-out by a scattered photon, rather than angle-induced dephasing of the spin wave. With a higher precision measurement, a delineation between these loss processes through a fit with multiple time constants may be possible, but this approach is not particularly convincing with the data at hand. Accounting for the technical losses in the setup by dividing out the transmission of a strong CW laser, which I take to be 30 %, and extrapolating to zero storage time yields a total internal efficiency of  $\eta_{\text{int}} = 4.6(9) \%$ . A measurement integrated for 20 min at 700 ns storage time yields a conditioned autocorrelation of the retrieved photons of  $g_{c,\text{ret}}^{(2)} = 0.503(93)$ . The lifetime as measured by the decrease in efficiency thus approximately corresponds to the crossing of  $g_{c,\text{ret}}^{(2)} > 0.5$  as well.

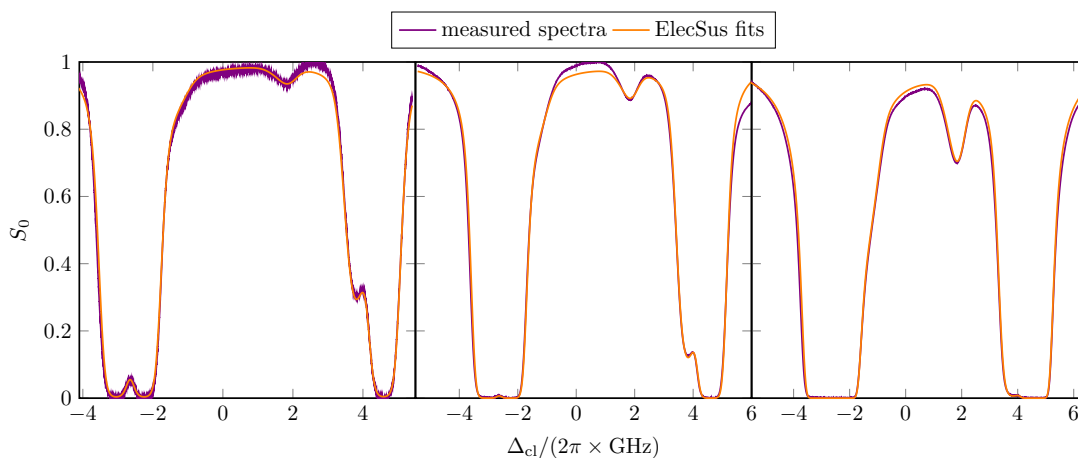
The photon statistics of the noise,  $g_{c,\text{noise}}^{(2)}$ , are not measured directly, as insufficient noise counts accumulate within a reasonable time to evaluate them meaningfully. As the control and pump lasers are not locked, an integration over night will not yield an accurate proxy for the experimental situation. Moreover, for all the reasons described in section 4.3, gaining any insights from measurements would have to involve modeling of the detector time response and temporal shape of the noise output. This seems somewhat excessive considering our additional knowledge about the nature of the noise. Sources of coherent noise for which we would have  $g_{c,\text{noise}}^{(2)} = 1$  are limited to leaked control laser light. Known possible thermal noise sources include uncorrelated SPDC photons from the source, which we have diligently eliminated by switching it off in time for the read out, as well as collisional fluorescence and four-wave mixing induced by the control. These latter scenarios can be distinguished by scanning the final filter etalon with the SPDC source blocked. I will show such measurements in the next section. They will reveal collisional fluorescence peaks on resonance with the natural oscillators  $F' \rightarrow F = 2$ , which easily dominate the noise at the signal frequency. No four-wave mixing peaks or control frequency contributions are visible after the filters. The expected value of  $g_{c,\text{noise}}^{(2)}$  is therefore 2 as non-thermal noise sources are well

excluded, so let us see whether this is compatible with the measured autocorrelation of the retrieved light. Using equation 2.7 to model an incoherent mixture of noise and signal, and setting  $g_{c, \text{noise}}^{(2)} = 2$ , yields an expected value of  $g_{c, \text{ret, theo}}^{(2)} = 0.205(29)$ . This agrees very well with the observed statistics, which is convincing evidence that this memory scheme eliminates the most problematic noise processes of the hyperfine scheme. Note, however, that this agreement is agnostic towards the noise statistics as setting  $g_{c, \text{noise}}^{(2)} = 1$  in equation 2.7 yields a result with well overlapped uncertainties. In other words, for incoherent admixtures of relatively small amounts of noise, its nature has little impact on the photon number statistics of the output.

To obtain these results an empirical optimization of parameters like control pulse power, width, and delay was performed as previously detailed. As the main results documenting the memory functionality are far less ambiguous here than they were in chapter 3, I will not spend a lot of space on this. Some characterizations, like the very broad scan of the detuning performed for the hyperfine memory (figure 3.12), could have been interesting but are not feasibly performed in a controlled manner i.e. all in one day. Frequency changes, in particular, are not implemented that rapidly on the side of the source, and also require retuning 5 etalons on the side of the memory. The time involved then demands a conscientious verification that other conditions remain comparable – adjustments beget adjustments. Moreover, the insights obtainable from the scan of a single parameter are few – there are too many interdependencies, i.e. in general only a local maximum in the parameter space is revealed. Practically, optimization across parameters therefore generally consists of a coarse scan to find the right range, followed by a finer scan to find the best performance, followed by weeks of minor adjustments and attempts to implement improvements before the suitability of the operation range is rechecked. The process is somewhat organic, and knowing the right time to attempt reoptimizing a setting calibrated in the past is a matter of experience. I thus make no claim to have found the perfect operational conditions, and refrain from showing misleading evidence to the contrary as that would only detract from the results. This does mean that the next steps with regards to parameter optimization will have to rely on guidance from a trustworthy simulation, which has since been implemented. We report on first insights gained from modeling in [382], and a more complete overview of our numerical approach is given in [383].

## 5.4 Further Characterization

In this section I have collected the most interesting “auxiliary” measurements characterizing the memory, beyond the storage experiments themselves. I will first detail how the atomic temperature is determined by spectroscopy, aided by ElecSus software from the Durham Quantum Light and Matter research section [415, 416]. I will then



**Figure 5.13:** *left (center) (right)* Spectra of the experiment vapor cell for thermistor measured heater temperatures of 63 °C (72 °C) (87 °C) and fits produced by ElecSus using the known cell parameters leaving only the temperature free. The resulting atomic temperatures are 52 °C (58.5 °C) (69.5 °C). The measured transmissions are normalized to a value far detuned from any atomic lines and corrected for power variation over the course of the laser scan to yield the Stokes parameter  $S_0$ .

quantify the quality of the state preparation at high atomic temperatures, checking how well the ground state is initialized and determining its lifetime. Finally I will show spectra of the noise before and after the spectral filters.

#### 5.4.1 Determining the Atomic Temperature

Temperature measurements using thermistors integrated in vapor cell heaters are not necessarily good proxies of the atomic temperature within the cell. As a gas will condense at the coldest point of a container, the atomic temperature is a function of the minimum vapor pressure anywhere in the vapor cell – usually in the cell stem. A thermistor measurement performed there would reflect the atomic temperature more accurately, however the best estimates are obtained through spectroscopy. Determining the atomic temperature to high accuracy this way is simplified by the excellent modeling and fitting implemented in ElecSus. To do so a weak probe laser is frequency scanned over the atomic lines of the test cell. This laser is simultaneously referenced with a Doppler-free spectroscopy to calibrate its absolute frequency, and with a power measurement to compensate for variation over the scanning region. Lee Weller describes the best practices in preparing spectral data for accurate modeling in his thesis [417], and I follow the advice given therein. The absolute frequency calibration by Doppler-free spectroscopy is included in the appendix 7.3.

Spectra obtained by scanning a weak probe with  $I \approx I_{\text{sat}}/150$  over the vapor cell



used in the Zeeman memory experiment for different settings of the cell heater are shown in figure 5.13. The known vapor cell parameters entering into these fits are a cell length of 75 mm as well as the frequency shift  $\Delta\nu_s = -37$  MHz and broadening  $\delta\nu_b = 82$  MHz induced by 5 Torr of nitrogen buffer gas. The  $^{85}\text{Rb}$  content of the cell is only specified as  $< 1\%$ . Adjusting this value manually to reproduce the measured dip in the middle of each spectrum yields a best estimate of 0.8%, which I then use consistently. A similar measurement performed with the vapor cell from the hyperfine memory at the same heater temperature (75 °C) used for the storage experiments in chapter 3 produces an atomic temperature of 63(1) °C. The statistical accuracy of ElecSus fitting is reported as 0.1 °C in [415], but naturally depends on the quality of the data provided. I hesitate to claim that I have achieved this level of shot-to-shot reproducibility, in particular as the normalization and correction for power variation during the scan does not always produce perfect results – see the upper middle and right flank of the central plot in figure 5.13 for an illustration of the kinds of discrepancies this causes. The uncertainty of 1 °C I give for the atomic temperatures in storage experiments is also intended to capture long term variations in what is produced by identical heater settings, as I do not remeasure the atomic temperature every day. Further note that it takes at least 12 h for the cell to be properly thermalized after a significant change in heater settings.

## 5.4.2 State Preparation

In dense atomic ensembles it is not self-evident that optical pumping effectively prepares the desired atomic state. The effect to blame for this is radiation trapping [418]. If the photon spontaneously emitted by an excited atom through radiative decay is sufficiently likely to be reabsorbed by another atom before it leaves the ensemble, then this “trapped” light produces a resonant background that competes with the pumping and degrades the achievable atomic polarization [194]. A well chosen buffer gas can extend the range of densities that can still be pumped by providing the excited state non-radiative decay pathways. This is called *quenching*. Molecular nitrogen, for instance, can relax excited Rb through collisions that transfer the energy to its vibrational states [419]. An empirical model for optical pumping in Rb vapor cells with  $< 50$  Torr buffer gas is described in [177]. A comparative analysis of obtainable polarizations in Cs with and without quenching in the context of memory experiments has also been published [377].

The Rb densities at which [177] cease being able to produce atomic polarizations near 100% with a little bit of nitrogen buffer gas are around  $n > 3 \times 10^{12} \text{ cm}^{-3}$ . Fortunately this is about an order of magnitude higher than the densities I have been using. Here the long vapor cell is helpful as the axial optical depth is already quite

#### 5.4. Further Characterization

---

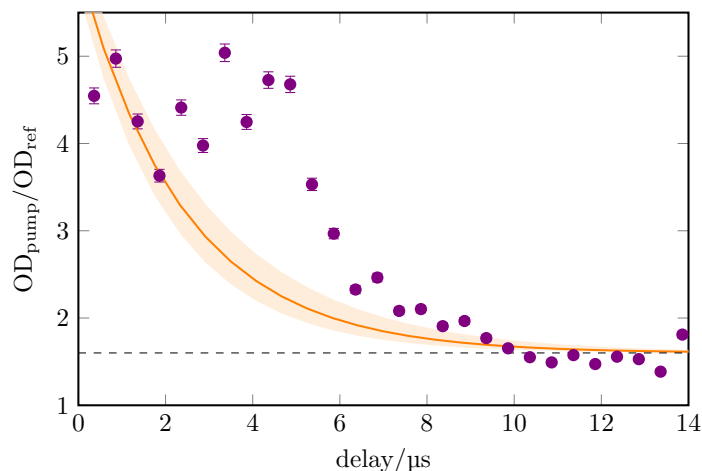
high for the density at  $T_{\text{at}} = 50^\circ\text{C}$ , which is just barely over  $n = 1 \times 10^{11} \text{ cm}^{-3}$ .<sup>9</sup> Nevertheless it is better to be sure, and effective state preparation not literally being impossible does not quite imply that it is also achieved. To characterize the polarization of an atomic ensemble over time a pump-probe experiment with relaxation in the dark can be preformed. The atoms are pumped for a while, then the pumping beams are switched off. The transmission of a weak probe through the vapor cell is measured as a function of the time after which the pumping ceases. Generally speaking, better accuracy is obtained when the prepared state is dark for the probe – i.e. when higher transmission corresponds to more atomic polarization. In the scenario of preparing the state  $|F = 2, m_F = 2\rangle$  this would require probing the hyperfine state preparation and the Zeeman state preparation separately, with probes of two frequencies. For this reason I opted to eschew the conventional wisdom and chose to probe the atomic polarization with  $\sigma^-$  polarized light on  $|F = 2\rangle \rightarrow |F' = 1\rangle$ .<sup>10</sup> This transition has the advantage that absorption increases monotonically with the atomic polarization. Assuming equal occupation without pumping, the expected increase in optical depth for perfect hyperfine pumping to  $F = 2$  is a factor  $8/5$ . The additional increase for perfect Zeeman state preparation can be calculated from the relative line strengths. The ratios of the strengths of the allowed transitions  $|F = 2, m_F = 2, 1, 0\rangle \rightarrow |F' = 1, m'_F = 1, 0, -1\rangle$  are  $6 : 3 : 1$  [366]. This yields an additional factor 3 for perfect preparation of  $|F = 2, m_F = 2\rangle$  over just  $|F = 2\rangle$ . The total ratio between the optical depths of the fully polarized and unprepared case is therefore  $4.8 : 1$ .

To check the extremes of the possible operating conditions of the memory I turn the vapor cell heater all the way up reaching an atomic temperature of  $70^\circ\text{C}$ . This corresponds to a density of  $n \approx 6 \times 10^{11} \text{ cm}^{-3}$ , so the atoms still are not quite in the [177] danger zone. Moreover, to test the pumping in a situation mimicking a memory experiment as closely as possible, I use a short pulse with a duration similar to that of the photons as a weak probe. The peak power of the probe is about 100 nW. As the optical depth is now very high, I also detune the probe from  $|F = 2\rangle \rightarrow |F' = 1\rangle$  by  $\Delta = -2\pi \times 1.1 \text{ GHz}$ . For a discussion on how the optical depths measured with detuned probes relate to the resonant optical depth in hot vapors with a buffer gas present see section 5.1.1. (The conversion is not relevant to calculating the ratio of optical depths with and without pumping as the detuning is held constant, but matters for determining the contribution to the absorption by the transition to the  $F' = 2$  excited state.) The atoms are pumped under conditions similar to those during storage experiments. The pumping power is about 18 mW and the repumping power is about 10 mW. An angle of 2.95(15) mrad is present between the spatial modes of pump and

---

<sup>9</sup>I calculate the density from the temperature via the liquid Rb vapor pressure as in [407].

<sup>10</sup>My thinking was that this would simplify the accounting. It is probably fair to say that I was wrong.



**Figure 5.14:** The increase in optical depth on the  $\sigma^-$  transition  $|F = 2\rangle \rightarrow |F' = 1\rangle$  due to optical pumping to  $|F = 2, m_F = 2\rangle$ . The minimum delay is 360 ns. The dashed line represents the OD increasing by a factor 8/5 due to the hyperfine state preparation, which works reliably and is independently measured to decay with an exponential lifetime of 1.4(2) ms. Data in the range 3  $\mu$ s–6  $\mu$ s do not follow the exponential decay due to a reproducible oscillation of unknown origin and are excluded from the fit. The fit yields a lifetime of 2.4(3)  $\mu$ s, which practically captures the speed of the decay reasonably well but not when it occurs, and an initial increase in OD by a factor 5.5(3). Error bars on the data represent the reproducibility of the data evaluation for reasonable variations of the pulse area integration region. The shaded region represents the 95 % confidence interval of the fit.

probe. The pumping duty cycle is fixed at 80%. After pumping the probe transmitted upon a variable delay is measured with a sensitive avalanche diode (MenloSystems APD210) and the pulse area is compared to what passes through the atoms without pumping. Figure 5.14 shows the relative optical depth as a function of the delay. While there is some scatter in the data stemming from the division by small numbers, the large discrepancy to the exponential fit in the range 3  $\mu$ s–6  $\mu$ s is not a sign of poor data quality. Instead it is the result of a widely reproducible oscillation. On longer time scales I find this oscillation has a period of about 20  $\mu$ s – the effect on short time scales qualitatively looks like a misalignment of the first oscillatory maximum from zero delay. So far I have not traced its origin, but have also observed it while the vapor cell heater (which heats with a current) is turned off, and also when the pumping beams are exactly on axis. Due to the large detuning, the off-resonant absorption by  $|F = 2\rangle \rightarrow |F' = 2\rangle$  needs to be accounted for when evaluating the initial increase in OD as well. Unfortunately the scale of its contribution varies with the quality of the Zeeman state preparation. If the quality is high its contribution to the optical depth

#### 5.4. Further Characterization

---


is 10%. While its contribution increases with worse Zeeman state preparation, that of  $|F = 2\rangle \rightarrow |F' = 1\rangle$  decreases faster, so this would not be a viable explanation for higher relative changes in the OD. Correcting, therefore, by this factor of 10% yields an initial increase in relative OD by 5.0(3), compatible only with a very high degree of atomic polarization.

The magnetic shield housing the vapor cell (Twinleaf MS-1L) has integrated coils for producing magnetic fields and gradients. I spent a short amount of time seeing whether I could influence the state preparation by applying fields. For “large” fields (relative to the coil capabilities, actually only about 5  $\mu\text{T}$ ) applied along the optical axis I measured  $\text{OD}_{\text{pump}}/\text{OD}_{\text{ref}} > 3$  quite consistently for delays up to 90  $\mu\text{s}$ . Unfortunately, the initial value was not as good as without field, and considering that only very short timescales are currently relevant to the memory purely due to atomic motion I decided to defer further tests.

In the course of characterizing the state preparation for another, related experiment in a different vapor cell Roberto Mottola and I considerably improved our approach. We now feel confident in determining atomic polarizations with uncertainties around 0.5%, and have achieved agreement to that degree in consistency-checking measurements probing different ground states. He reports on these experiments in his thesis [383], and with hindsight I would not have gone about the measurement as I have described above. Implementing this method in this setup will enable both a more accurate determination of the atomic polarization and a more thorough optimization of the state preparation in future experiments, and will be a key tool towards extending its lifetime to enable long storage times.

##### 5.4.3 Noise Spectra

Spectral information about the noise, both directly after the memory and after the spectral filters, can provide additional insights. The former reveals what physical processes produce additional photons in the vapor cell. It is, for instance, possible to distinguish collisional fluorescence from four-wave mixing even with a rather coarse frequency resolution, as the former always occurs at the natural oscillator frequency while peaks due to the latter vary their frequency with  $\Delta$ . Meanwhile, the latter delivers reasonably quantitative evidence on whether sufficient spectral filtering is used, or whether the noise performance can be improved with additional etalons. To obtain this information I run memory experiments without input, then temperature scan an etalon filtering the memory emissions before measuring them as in a memory experiment with the single photon detectors. The triggering logic used can be sourced from the detection of SPDC idler photons, but I have not observed any significant differences between this and using deterministic triggers. The resonant frequencies of fused silica etalons



**Figure 5.15:** *left* A spectrum of the noise attenuated but without spectral filtration, operating the memory at  $\Delta = 2\pi \times 1.7$  GHz or 900 MHz to the blue of  $F' = 2$ . The orange curve demonstrates the calibration of the frequency axis by a second measurement including a small amount of light at the signal frequency. Its contribution is labeled  $\mathcal{E}$ . *right* A spectrum of the noise of the memory operating at  $\Delta = -2\pi \times 600$  MHz, after filtration by two 4 mm etalons statically set to transmit the signal. A peak labeled  $\Omega$  remains visible at the control frequency, but at the signal frequency only noise originating from collisional fluorescence remains. As this noise source is spectrally near the signal, occurring at the natural transition frequency, filtration by further etalons is subject to rapidly diminishing returns. The trade-off of slightly lower noise vs. reduced signal transmission and possibly bandwidth restriction seems costly, although a direct test may be worthwhile.

vary by  $-2.391$  GHz  $\text{K}^{-1}$  at 795 nm, so it is important to scan sufficiently slowly. I use scan speeds of  $0.05$  K  $\text{s}^{-1}$ , or  $-120$  MHz  $\text{s}^{-1}$  equivalently, in all measurements. A scan should also be sufficiently broad, preferably covering multiple free spectral ranges so that a central one can be selected for maximal scan linearity, and to calibrate the scale of the frequency axis if necessary. Ramping over  $30^\circ\text{C}$  and picking out the data from the middle of the scan range I have generally been able to omit this step. In particular, the data that follow are presented with the frequency axis scaled purely as is expected from thermal expansion. There is some variation between etalons in how smoothly they expand. If I were to speculate, I would blame this on random differences in material stresses from their mounts. For thin etalons it suffices to center the interesting features in the scan range. To have an absolute frequency reference I perform these scans twice, adding a small amount of light at the signal frequency the second time round. For well chosen intensities of this reference this produces two nearly identical spectra, but gives one an unmistakable additional peak at a known absolute frequency. The response of the etalons is smoother and the resulting spectra are cleaner when they are allowed to expand rather than contract, so for optimal data quality I always ramp towards higher temperatures.

#### 5.4. Further Characterization

---

Spectra of the noise before and after filtration for 2 different scenarios are shown in figure 5.15. A spectrum directly after the memory is shown on the left. For the purpose of illustration I am showing one collected operating at about  $\Delta = 2\pi \times 1.7$  GHz, or 900 MHz to the blue of  $F' = 2$ . The atomic temperature here is  $T_{\text{at}} = 59(1)$  °C. The light emitted by the memory is attenuated by neutral density filters before it is detected to avoid saturation. In this plot I have also included the spectrum taken with a small amount of light at the signal frequency for calibration to show what is going on. A 2 mm etalon is scanned, yielding a minimum feature width of about 1.2 GHz due to the etalon bandwidth. There is some variation in how accurate the frequency axes turn out in these measurements. For greater confidence in labeling the peaks it is often worth checking directly whether the peaks move with the control frequency. The features at the natural transitions frequencies are labeled CF1 and CF2 respectively, corresponding to the ground state the relaxing atom ends up in. Their frequencies line up with expectations nicely. The small feature I have labeled FWM falls shy of  $2 \times \Delta_{\text{hf}}$ , but as its frequency follows the control frequency the identification does not seem too bold a claim to make. The very steep feature whose full amplitude I have cut off is the result of an EOM re-locking during the measurement, briefly transmitting a large amount of control laser light. The final state of the unfiltered memory emissions under the conditions of section 5.3 are likely considerably better, both due to a somewhat lower temperature and later optimizations performed directly on the memory performance.

In the spectrum shown on the right hand side of figure 5.15, the memory emissions are first filtered by two 4 mm etalons whereupon the third and last 4 mm etalon of the memory filtering stage is scanned. The operation point is  $\Delta = -2\pi \times 600$  MHz and the atomic temperature is  $T_{\text{at}} = 52(1)$  °C. Thanks to some luck with the etalon's temperature scanning behavior the frequency axis matches expectations unusually accurately. The dashed lines on the fluorescence peaks correspond to the transitions  $|F' = 2\rangle \rightarrow |F = 2, 1\rangle$ . The etalon bandwidth of 550 MHz would allow for the resolution of a separate fluorescence peak from  $F' = 1$ , which is notably absent even when the data are plotted logarithmically. This could be interpreted as the state preparation inducing sufficient polarization in both ground states that the lasers' ability to address  $F' = 1$  is suppressed even when they do find an atom able to scatter light. For further comparison, a more detailed spectrum of the kind of additional structure visible in the hyperfine memory scheme acquired by this method is shown in figure 7.4. Note that the count rates measured here are not directly comparable to those measured in the retrieval region during storage experiments as the data are not time gated. The correct comparison would be an integrated rate. To isolate the spectrum of light produced only during read out a similar but much more time consuming approach is required. The absolute noise rates within the read out region are so low (about  $20 \text{ s}^{-1}$ ), that to accumulate sufficient counts for a meaningful contrast a step-wise scan of the etalon

and lengthy integration of the memory output at every etalon temperature would be needed. Further, a hardware based gate, e.g. a fiber switch, should be used to avoid storing and analyzing enormous amounts of worthless timestamp data outside of the region of interest – consider the small duty cycle. Such a time resolved spectral measurement would truly be the last word on memory noise.

In summary, these spectra add the final piece to the convergence of evidence concerning the nature of noise present in the Zeeman memory. The complete absence of features bearing the characteristics of FWM provides sound reason to conclude that noise is dominated by collisional fluorescence. The visibility of a spectral peak at the control frequency after two etalons proves that our filtering is not excessive, while evidence that it is sufficient is provided in section 5.2.2. To eliminate the remaining noise two approaches are conceivable. Optimally, the quality and longevity of the ground state preparation should be further improved. The best way to deal with noise is always to avoid creating it, and atoms therein cannot be excited to fluoresce. Beyond that, spectral filtration of collisional fluorescence is possible when  $\Delta$  is sufficiently large to separate the signal from its broad lines at the natural transitions frequencies. As the memory scheme itself limits how far from the lines the lambda system may be detuned, let alone the technical difficulty of providing sufficient control intensity at greater  $\Delta$ , an attempt at this should only be made upon further improving the noise measurements to allow for fast, time-gated spectral analysis. Systematic gains in this regard are conceivable, but must be weighed carefully against efficiency trade-offs when the goal is to further improve the state accuracy of retrieved photons.





## Chapter 6

# Outlook

*Science never sleeps, so I'm singing it a lullabye,  
Of wolves in sheep's clothing – devils in disguise,  
And lab-coats draped over big black lies.*

Martin Walkyier

---

In this thesis I have described multiple iterations of a hot vapor quantum memory and a spontaneous parametric downconversion photon-pair source operating in a technologically relevant bandwidth regime. My report culminates in their interfacing and a successful demonstration of single photon storage. An obvious next step is one further characterization, namely evaluating the fidelity of the interface with Hong-Ou-Mandel interference. This is of course interesting in its own right, and could simultaneously be considered a small scale implementation of a heralded photon source synchronization. Interference could be measured between a stored photon and one emitted by the source directly, or between subsequently stored photons. Beyond the interferometer, this also requires a small revision of the electronics to enable triggered storing and retrieving. This change should be trivial as timing demands are not affected – all that is required is a flip-flop.

Simulation guided optimization of the operating regime may then proceed to yield better functionality, or unambiguously identify remaining limitations. A conclusion near at hand and already confirmed by our model taking into account the transverse mode profiles of the signal and control is that greater homogeneity of the control Rabi frequency would significantly improve the efficiency. A control with a flat-top intensity profile corresponding to the ideal Rabi frequency may be a little utopian considering the required  $\Omega$ , but a Gaussian control mode with a waist just twice as large as the signal's would already yield considerable gains. The intensity requirements of this approach do not scale as poorly as might be expected, as the required peak Rabi

---

frequency to optimize storage decreases with greater spatial homogeneity. For more information on these models refer to [382] and [383].

Further investigation into improving the state preparation using magnetic fields could also be fruitful, as such methods have yielded very long state lifetimes elsewhere [197]. Longer lived state preparation could obviate or at least reduce the need for pumping between storage attempts, allowing the memory to operate at higher rates theoretically – when not limited by source rates. With the maximum currently available control power, a control mode with a 1 mm diameter could still achieve the same  $\Omega$  used to obtain the results of chapter 5. Admittedly, operating the control generation under the conditions producing peak power currently induces undesirable side-effects, particularly worse EOM performance and correspondingly more noise. Alternative means of generating the control pulses, taking inspiration from our SOA switches and pulsing the current to a TA seeded in CW, have shown a great deal of initial promise and will soon be tested in the context of memory experiments. Optimal control in a mode of this size puts the memory well within range of vapor cell fabrication techniques matching the cell size to the control [369, 420]. Many orders of magnitude in lifetime could thus be gained relatively simply by eliminating motional loss, assuming the cell walls could be coated against spin destroying collisions. Finally, with recently obtained superconducting nanowire single photon detectors, the interface will experience a jitter reduction that may improve the figures of merit slightly. These detectors will also allow for more precise specification of the figures of merit, as they have smaller uncertainties on their detection efficiencies.

Storing and retrieving single photons in our interfaced system was a difficult problem. Nevertheless, it is only the beginning of interesting experiments that are opened up by this platform. These include a plethora of elementary demonstrations of protocols described in section 2.3, but also further interfacing experiments with other systems. Based on the experiments of chapter 5, the expected signal to noise ratio interfacing the Zeeman memory with a source of Rb wavelength single photons with efficiency  $\eta$ , and no added noise, is  $\eta \times \frac{\text{SNR}_{\text{mem}}}{\eta_h} = \eta \times 37(6)$ . This indicates that the memory would be compatible with single photons extracted from GaAs/AlGaAs quantum dots – approximately at the level that the hyperfine memory was compatible with laser pulses – for extraction efficiencies of only 10 %, which is still in the range of what is possible with comparatively simple methods like a solid immersion lens [421]. Such quantum dot single-photon sources are investigated by our collaborators in the laboratories of Richard Warburton [269, 270, 280]. Moreover, the parallel development in our lab of a memory implementing a lambda scheme isolated by applying a high magnetic field is almost complete. An interface of these systems could also be envisioned, implementing a kind of memory bus. In the longer term, an iteration of the combination of a single photon source and quantum memory operating at high rates and bandwidth, in a simple

experimental setup as demonstrated here, could form the basis for QIP applications. As further commented upon at the end of section [2.3.4](#), the synchronization of photons and gates in local networks is a particularly suited task for fast and high bandwidth memories.



# Chapter 7

## Appendices

*So here's some more words*

*And non sequiturs*

*Because blah blah blah blah*


*Is all that you heard*

Al Jourgensen


---

### 7.1 Full Level Diagram for Hyperfine Memory

This level diagram provides the full context for the prospect of photon storage in a hot vapor memory scheme based on degenerate hyperfine levels, as discussed in chapter 3. The important consideration in regard to the principle viability is that for states linked by lambda systems the net result will be interference, for which an optimal detuning may be found and which may be compensated for with greater optical depth. The problematic state for storage, marked with the dotted line in figure 7.1, is  $|F' = 2, m'_F = 0\rangle$  because it can absorb the signal interference free, lacking any control coupling. This excludes a large region of possible detunings near the resonance of this line from consideration due to the competing absorption process, in particular its presence is modeled to reduce the memory efficiency by at least half in a fairly symmetric region of about 1 GHz around the line due to photon absorption without storage under our experimental conditions. In the results described in section 3.4, where  $|F' = 2, m'_F = 0\rangle$  is a good 1.7 GHz blue detuned from the memory working point, its impact on the storage efficiency is on the order of 10%.



**Figure 7.1:** This is a complete level diagram of the transitions allowed to orthogonal linear polarizations on the different hyperfine transitions of  $^{87}\text{Rb}$  D<sub>1</sub>. The quantization axis is chosen such that the light incident on  $F=2 \rightarrow F'$  is  $\pi$ -polarized, leaving the orthogonal linear polarization incident on  $F=1 \rightarrow F'$  in the representation of the equal superposition  $\sigma^+ + \sigma^-$ . Tables of the transition strengths and their signs can be found online [115].



**Figure 7.2:** Measured detector response of two SPADs in HBT configuration as used for roughly half of this thesis. As they are quite tight, the 95 % confidence intervals of the fits are directly used as the line thickness, down to a fixed minimum width. The FWHM of the response is estimated at 1.09(5) ns by the sech fit and at 1.21(5) ns by the Gaussian. The uncertainties include the fit confidence and one 30 ps bin width as an estimate of systematic error.

## 7.2 HBT Detection Jitter Measurement

The time resolution of the single photon detectors used in these experiments (Excelitas, SPCM-AQRH-16) is only specified as  $\delta\tau > 350$  ps. Two detectors arranged in Hanbury Brown and Twiss configuration are thus expected to achieve a minimum detection jitter of  $\sqrt{2}\delta\tau > 500$  ps. As I discuss in section 5.2.1, it is ambiguous whether this number is intended to refer to one standard deviation in the delay of the output voltage pulse or the full width half maximum of the detector's impulse response. This poses a problem, as the models shown in figure 4.3 require accounting for the response function of the HBT to accurately reproduce the data. Unfortunately, a specification stated as an inequality, without data from which we could derive exactly what (if anything) was measured by the manufacturer, is difficult to enter into a model. It was therefore necessary to determine the detector jitter empirically. I know now that for the detectors as they were coupled to gather the time-tagged data of chapter 3 and section 4.3 we probably did not succeed in optimizing the detectors' timing accuracy. For instance, only one lens was used to focus the beam on both detectors.

To measure the instrument (or impulse) response function of an optical detector a short optical pulse is required. Any pulse shape measured in time consists of a convolution of the real signal and the detector response. If the signal is temporally much shorter than the detector's resolution, then it can be approximated as a delta-distribution. As these are the identity element of convolution, the measured shape is then equal to the detector response. To this end, my colleagues Alisa Javadi and Natasha Tomm kindly let me use their picosecond laser (Coherent Mira 900) and fast time tagger (Swabian Instruments). The result of sending 12 ps pulses from this laser onto the HBT is shown in figure 7.2. The determined line shape directly enters the models of figure 4.3. After a peak dominated by fast transients, the temporal response of SPADs is exponential [422]. While the FWHM of the response depends only weakly on what function is used to model it, a hyperbolic secant therefore does a better job of capturing the tails than a Gaussian. After this measurement the HBT was realigned with more degrees of freedom, and I expect the jitter since then to be near the minimum achievable amount.


## 7.3 Absolute Frequency Calibration by Rb Spectroscopy

To calibrate the absolute frequency of a scanned laser without assuming the scan is perfectly linear, all the transitions visible in a Doppler-free natural Rb spectroscopy can be used. The transitions that are reliably sufficiently prominent to be identified in the spectra I measure are listed in table 7.1, and the conversion of the scan-time axis recorded by the oscilloscope into an absolute frequency axis is illustrated in figure 7.3.

### 7.3. Absolute Frequency Calibration by Rb Spectroscopy

Transition	$\Delta_{\text{cl}}, 2\pi \times \text{GHz}$
$^{87}\text{Rb } F = 2 \rightarrow F' = 1$	-3.016
$^{87}\text{Rb } F = 2 \rightarrow F' = 2$	-2.201
$^{85}\text{Rb } F = 3 \rightarrow F' = 2$	-1.498
$^{85}\text{Rb } F = 3 \rightarrow F' \text{ crossover}$	-1.317
$^{85}\text{Rb } F = 3 \rightarrow F' = 3$	-1.136
$^{85}\text{Rb } F = 2 \rightarrow F' \text{ crossover}$	1.719
$^{85}\text{Rb } F = 2 \rightarrow F' = 3$	1.900
$^{87}\text{Rb } F = 1 \rightarrow F' = 1$	3.819
$^{87}\text{Rb } F = 1 \rightarrow F' = 2$	4.634

**Table 7.1:** Table of transition frequencies in a natural Rb reference spectroscopy relative to the ElecSus definition of the Rb center line.




**Figure 7.3:** The absolute frequency of a scanned laser is calibrated by referencing a Doppler-free spectroscopy of natural rubidium. The  $y$ -axis of the spectrum is the normalized and power-variation corrected transmission through the reference cell. Nine transitions are identified and a (minimally) quadratic fit to their known frequencies calibrates the absolute frequency of the laser.



Here the reference point is the Rb *center line frequency* definition used in ElecSus. This definition is  $\nu_{\text{cl}} = 377.107\,407$  THz, which is an average of the center line frequencies of each of the Rb isotopes, collected for example in [115], weighted by their natural abundances. This frequency axis can then be applied to another, simultaneously recorded spectrum. This enables, for example, fitting the atomic temperature of the experiment vapor cell by ElecSus, as described in section 5.4, and bandwidth measurements of the filtering etalons. To compensate for power variation over the scan range a third trace monitoring just the power can be recorded and used to correct a slanted spectrum. I have semi-automated this whole process with a MATLAB script requiring only a few mouse clicks to roughly point out the Doppler-free peaks.

## 7.4 Hyperfine Memory Noise Spectrum

The same technique for measuring noise spectra presented in section 5.4.3 is used to measure the noise of a memory using the hyperfine storage scheme. The experimental setup used is an intermediate hybrid of the ones described in chapters 3 and 5, in particular a dedicated pumping laser switched by a SOA is already in use. The memory is run without input and a single 2 mm etalon filtering the output is slowly temperature scanned. A second measurement with a weak input at a known frequency calibrates the axis. The results for two different detunings,  $\Delta_{\text{red}} = -2\pi \times 600$  MHz and  $\Delta_{\text{blue}} = 2\pi \times 1.3$  GHz, both relative to  $F = 2 \rightarrow F' = 1$  are shown in figure 7.4 (recall that in the hyperfine scheme this is the transition nearest to the control frequency). As considerable structure is visible, the data are plotted on an absolute frequency axis also relative to the atomic transition  $F = 2 \rightarrow F' = 1$ . This allows for unambiguous distinction of the spectral peaks. The two peaks that are well aligned in the two traces and are marked with dashed lines correspond to fluorescence on the lines  $F = 2 \rightarrow F' = 1$  and  $F = 1 \rightarrow F'$ . The strong peak near zero is the control laser itself. The remaining structure can presumably be attributed to various FWM processes, although no attempt is made to identify it further.



**Figure 7.4:** Noise spectra recorded in the hyperfine memory scheme at two operational points:  $\Delta_{\text{red}} = -2\pi \times 600 \text{ MHz}$  and  $\Delta_{\text{blue}} = 2\pi \times 1.3 \text{ GHz}$ , both relative to  $F = 2 \rightarrow F' = 1$ .

# List of Figures

1.1	The memory hierarchy in a typical classical computer . . . . .	3
1.2	Schematic of how interference experiments are performed with photons: Hanbury Brown and Twiss effect and Hong-Ou-Mandel interference . .	14
2.1	Variable delay line (optical loop) as a simple memory . . . . .	19
2.2	A basic lambda scheme . . . . .	23
2.3	Three level scheme for CRIB-style photon echo memories . . . . .	26
2.4	Three level scheme for AFC-style photon echo memories . . . . .	28
2.5	Three level scheme for EIT memories and schematic plot of the suscepti- bility . . . . .	32
2.6	Level scheme illustrating the four wave mixing process . . . . .	36
2.7	Three level scheme for Raman memories . . . . .	40
3.1	Level diagram of a memory scheme using the hyperfine states of the Rb D <sub>1</sub> line . . . . .	68
3.2	Overview of the hyperfine memory setup components . . . . .	75
3.3	Hyperfine memory setup: signal pulse generation . . . . .	76
3.4	Characterization of the weak coherent signal pulse stored in the hyperfine memory . . . . .	78
3.5	Hyperfine memory setup: control pulse generation . . . . .	79
3.6	Hyperfine memory vapor cell photos . . . . .	81
3.7	Etalon design iterations . . . . .	84
3.8	Arrival time histograms for storage experiments in the hyperfine memory	88
3.9	Lifetime measurement of the hyperfine memory . . . . .	89
3.10	First- and second-order coherence measurement in the hyperfine memory	90
3.11	Further characterization of the hyperfine memory . . . . .	91
3.12	Data and simulations of memory efficiency versus detuning . . . . .	92
4.1	Setup of the triple resonant cavity SPDC source . . . . .	104
4.2	Spectra of the triple resonant cavity SPDC source . . . . .	106
4.3	Correlations measurements from the SPDC source . . . . .	110
4.4	Alterations to the source to improve interfaced noise performance . . .	115
5.1	Level diagram of a memory scheme using Zeeman states of the Rb D <sub>1</sub> line	120

5.2	Optical depth and relative line strengths of the two Rb $D_1$ excited states as functions of detuning . . . . .	123
5.3	Photo of the Zeeman memory . . . . .	124
5.4	Comprehensive sketch of the interfaced Zeeman memory and SPDC source	125
5.5	The effect of timing jitter on memory performance . . . . .	128
5.6	Characterization of the control pulse for the Zeeman memory: shape and spectral purity . . . . .	133
5.7	Characterization of the control pulse for the Zeeman memory: extinction	135
5.8	Performance of semiconductor optical amplifiers as optical switches . . .	137
5.9	Transmission through right-triangular calcite prisms (Fresnel) . . . . .	140
5.10	Arrival time histograms for storage experiments in the Zeeman memory: linear scale . . . . .	145
5.11	Arrival time histograms for storage experiments in the Zeeman memory: logarithmic scale . . . . .	145
5.12	Lifetime measurement of the Zeeman memory . . . . .	147
5.13	Determining the atomic temperature via spectroscopy . . . . .	150
5.14	Characterization of the Zeeman ground state preparation . . . . .	153
5.15	Noise spectra of the Zeeman memory before and after filtration . . . . .	155
7.1	Complete level diagram of transitions allowed to orthogonal linear polarizations on $^{87}\text{Rb } D_1$ . . . . .	164
7.2	Jitter measurement of the HBT configured photon detectors . . . . .	164
7.3	Absolute laser frequency calibration by Rb spectroscope . . . . .	166
7.4	Noise spectrum produced by the hyperfine memory scheme . . . . .	168

## List of Tables

3.1	Etalon parameters . . . . .	84
4.1	Monolithic ppKTP cavity parameters . . . . .	101
4.2	Table of double-resonant SPDC source parameters . . . . .	113
7.1	Transition frequencies in a natural Rb relative to the center line . . . . .	166

# Bibliography

- [1] T. W. Hänsch, *Repetitively pulsed tunable dye laser for high resolution spectroscopy*, [Applied Optics](#) **11**, 895 (1972).
- [2] T. W. Hänsch, *A passion for precision*, in K. Grandin (editor), *Les Prix Nobel (Nobel Foundation, Stockholm)* (2006).
- [3] G. E. Moore, *Cramming more components onto integrated circuits*, *Electronics* **38**, 114 (1965).
- [4] C. Obermair, F.-Q. Xie, and T. Schimmel, *The single-atom transistor: Perspectives for quantum electronics on the atomic-scale*, [Europhysics News](#) **41**, 25 (2010).
- [5] I. L. Markov, *Limits on fundamental limits to computation*, [Nature](#) **512**, 147 (2014).
- [6] H. J. Caulfield and S. Dolev, *Why future supercomputing requires optics*, [Nature Photonics](#) **4**, 261 (2010).
- [7] M. M. Shulaker, G. Hills, N. Patil, H. Wei, H.-Y. Chen, H.-S. P. Wong, and S. Mitra, *Carbon nanotube computer*, [Nature](#) **501**, 526 (2013).
- [8] G. Hills, C. Lau, A. Wright, S. Fuller, M. D. Bishop, T. Srimani, P. Kanhaiya, R. Ho, A. Amer, Y. Stein, D. Murphy, Arvind, A. Chandrakasan, and M. M. Shulaker, *Modern microprocessor built from complementary carbon nanotube transistors*, [Nature](#) **572**, 595 (2019).
- [9] D. A. Patterson and J. L. Hennessy, *Computer organization and design* (Elsevier LTD, Oxford) (2013).
- [10] K. Itoh, M. Horiguchi, and H. Tanaka (editors), *Ultra-low voltage nano-scale memories* (Springer, Boston, MA) (2007).
- [11] E. Conrad, S. Misenaar, and J. Feldman, *Domain 2: Asset security (protecting security of assets)*, in *CISSP Study Guide*, pp. 81–101 (Elsevier) (2016).

- [12] T. Kosub, M. Kopte, R. Hühne, P. Appel, B. Shields, P. Maletinsky, R. Hübner, M. O. Liedke, J. Fassbender, O. G. Schmidt, and D. Makarov, *Purely antiferromagnetic magnetoelectric random access memory*, [Nature Communications](#) **8** (2017).
- [13] D. Bossini, M. Pancaldi, L. Soumah, M. Basini, F. Mertens, M. Cinchetti, T. Satoh, O. Gomonay, and S. Bonetti, *Ultrafast amplification and nonlinear magnetoelastic coupling of coherent magnon modes in an antiferromagnet*, [Physical Review Letters](#) **127**, 077202 (2021).
- [14] J. Åkerman, *Toward a universal memory*, [Science](#) **308**, 508 (2005).
- [15] D. Reinsel, J. Gantz, and J. Rydning, *Data age 2025: The digitization of the world from edge to core*, techreport US44413318, International Data Corporation (2018).
- [16] P. Anderson, R. Black, A. Cerkauskaitė, A. Chatzieftheriou, J. Clegg, C. Dainty, R. Diaconu, R. Drevinskas, A. Donnelly, A. L. Gaunt, A. Georgiou, A. G. Diaz, P. G. Kazansky, D. Lara, S. Legtchenko, S. Nowozin, A. Ogus, D. Phillips, A. Rowstron, M. Sakakura, I. Stefanovici, B. Thomsen, L. Wang, H. Williams, and M. Yang, *Glass: A new media for a new era?*, in *10th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 18)* (USENIX Association, Boston, MA) (2018).
- [17] S. S. P. Parkin, M. Hayashi, and L. Thomas, *Magnetic domain-wall racetrack memory*, [Science](#) **320**, 190 (2008).
- [18] D. C. Gilmer, T. Rueckes, and L. Cleveland, *NRAM: A disruptive carbon-nanotube resistance-change memory*, [Nanotechnology](#) **29**, 134003 (2018).
- [19] M. Arita, M. Hagiya, M. Takinoue, and F. Tanaka, *DNA memory*, in *Handbook of Natural Computing*, pp. 1281–1318 (Springer Berlin Heidelberg) (2012).
- [20] L. Szilard, *Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen*, [Zeitschrift für Physik](#) **53**, 840 (1929).
- [21] C. E. Shannon, *A mathematical theory of communication*, [Bell System Technical Journal](#) **27**, 379 (1948).
- [22] M. M. Wilde, *Quantum information theory* (Cambridge University Press) (2018).
- [23] S. Höst, *Information and communication theory*, chapter Information measures for continuous variables, pp. 213–235 (John Wiley & Sons) (2019).
- [24] C. E. Shannon, *The redundancy of English*, in C. Pias (editor), *Cybernetics The Macy Conferences 1946-1953 - The Complete Transactions*, volume 1, pp. 248–272 (University of Chicago Press) (2016).

- [25] J. W. Jeong, Y.-E. Choi, W.-S. Kim, J.-H. Park, S. Kim, S. Shin, K. Lee, J. Chang, S.-J. Kim, and K. R. Kim, *Tunnelling-based ternary metal–oxide–semiconductor technology*, [Nature Electronics](#) **2**, 307 (2019).
- [26] Y.-H. Luo, H.-S. Zhong, M. Erhard, X.-L. Wang, L.-C. Peng, M. Krenn, X. Jiang, L. Li, N.-L. Liu, C.-Y. Lu, A. Zeilinger, and J.-W. Pan, *Quantum teleportation in high dimensions*, [Physical Review Letters](#) **123** (2019).
- [27] A. S. Holevo, *Bounds for the quantity of information transmitted by a quantum communication channel*, [Problemy Peredachi Informatsii](#) **9**, 3 (1973).
- [28] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge University Press) (2010).
- [29] D. Gottesman, *The Heisenberg representation of quantum computers*, in *22nd International Colloquium on Group Theoretical Methods in Physics*, pp. 32–43 (1998).
- [30] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Quantum entanglement*, [Reviews of Modern Physics](#) **81**, 865 (2009).
- [31] E. Chitambar and G. Gour, *Quantum resource theories*, [Reviews of Modern Physics](#) **91**, 025001 (2019).
- [32] L. Masanes, *All bipartite entangled states are useful for information processing*, [Physical Review Letters](#) **96** (2006).
- [33] M. Horodecki, P. Horodecki, and R. Horodecki, *Separability of mixed states: Necessary and sufficient conditions*, [Physics Letters A](#) **223**, 1 (1996).
- [34] S. Gharibian, *Strong NP-hardness of the quantum separability problem*, *Quantum Information and Computation* **10**, 343 (2010).
- [35] M. Lewenstein, B. Kraus, J. I. Cirac, and P. Horodecki, *Optimization of entanglement witnesses*, [Physical Review A](#) **62** (2000).
- [36] M. B. Plenio and S. Virmani, *An introduction to entanglement measures*, *Quantum Information and Computation* **7**, 1 (2007).
- [37] W. K. Wootters and W. H. Zurek, *A single quantum cannot be cloned*, [Nature](#) **299**, 802 (1982).
- [38] B. M. Terhal, *Quantum error correction for quantum memories*, [Reviews of Modern Physics](#) **87**, 307 (2015).
- [39] D. P. DiVincenzo, *The physical implementation of quantum computation*, [Fortschritte der Physik](#) **48**, 771 (2000).
- [40] D. Loss and E. V. Sukhorukov, *Probing entanglement and nonlocality of electrons in a double-dot via transport and noise*, [Physical Review Letters](#) **84**, 1035 (2000).

- [41] A. Bienfait, K. J. Satzinger, Y. P. Zhong, H.-S. Chang, M.-H. Chou, C. R. Conner, É. Dumur, J. Grebel, G. A. Peairs, R. G. Povey, and A. N. Cleland, *Phonon-mediated quantum state transfer and remote qubit entanglement*, [Science](#) **364**, 368 (2019).
- [42] A. Bienfait, Y. P. Zhong, H.-S. Chang, M.-H. Chou, C. R. Conner, É. Dumur, J. Grebel, G. A. Peairs, R. G. Povey, K. J. Satzinger, and A. N. Cleland, *Quantum erasure using entangled surface acoustic phonons*, [Physical Review X](#) **10**, 021055 (2020).
- [43] T. E. Northup and R. Blatt, *Quantum information transfer using photons*, [Nature Photonics](#) **8**, 356 (2014).
- [44] E. Knill, R. Laflamme, and G. J. Milburn, *A scheme for efficient quantum computation with linear optics*, [Nature](#) **409**, 46 (2001).
- [45] R. Raussendorf and H. J. Briegel, *A one-way quantum computer*, [Physical Review Letters](#) **86**, 5188 (2001).
- [46] G. J. Milburn, *Quantum optical Fredkin gate*, [Physical Review Letters](#) **62**, 2124 (1989).
- [47] I. L. Chuang and Y. Yamamoto, *Simple quantum computer*, [Physical Review A](#) **52**, 3489 (1995).
- [48] L.-M. Liang and C.-Z. Li, *Realization of quantum SWAP gate between flying and stationary qubits*, [Physical Review A](#) **72** (2005).
- [49] Y.-A. Chen, S. Chen, Z.-S. Yuan, B. Zhao, C.-S. Chuu, J. Schmiedmayer, and J.-W. Pan, *Memory-built-in quantum teleportation with photonic and atomic qubits*, [Nature Physics](#) **4**, 103 (2008).
- [50] 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).
- [51] 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** (2007).
- [52] 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).
- [53] H. J. Kimble, *The quantum internet*, [Nature](#) **453**, 1023 (2008).
- [54] S. Wehner, D. Elkouss, and R. Hanson, *Quantum internet: A vision for the road ahead*, [Science](#) **362**, eaam9288 (2018).



- [55] 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).
- [56] J. L. Dodd and M. A. Nielsen, *Simple operational interpretation of the fidelity of mixed states*, [Physical Review A](#) **66** (2002).
- [57] C. Mewes and M. Fleischhauer, *Decoherence in collective quantum memories for photons*, [Physical Review A](#) **72**, 022327 (2005).
- [58] M. Razavi, M. Piani, and N. Lütkenhaus, *Quantum repeaters with imperfect memories: Cost and scalability*, [Physical Review A](#) **80**, 032301 (2009).
- [59] 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).
- [60] 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).
- [61] 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).
- [62] 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).
- [63] D. V. Vasilyev, I. V. Sokolov, and E. S. Polzik, *Quantum memory for images: A quantum hologram*, [Physical Review A](#) **77**, 020302 (2008).
- [64] O. Firstenberg, M. Shuker, A. Ron, and N. Davidson, *Colloquium: Coherent diffusion of polaritons in atomic media*, [Reviews of Modern Physics](#) **85**, 941 (2013).
- [65] C. C. Gerry and P. Knight, *Introductory quantum optics*, (Cambridge University Press, Cambridge) (2004).
- [66] R. Loudon, *The quantum theory of light* (OUP Oxford) (2000).
- [67] 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).

- [68] R. Ikuta, T. Kobayashi, K. Matsuki, S. Miki, T. Yamashita, H. Terai, T. Yamamoto, M. Koashi, T. Mukai, and N. Imoto, *Heralded single excitation of atomic ensemble via solid-state-based telecom photon detection*, [Optica](#) **3**, 1279 (2016).
- [69] A. M. Brańczyk, *Hong-Ou-Mandel interference*, [arXiv:1711.00080](#) (2017).
- [70] B. Kambs and C. Becher, *Limitations on the indistinguishability of photons from remote solid state sources*, [New Journal of Physics](#) **20**, 115003 (2018).
- [71] W. Zhi, R. Guobin, L. Shuqin, and J. Shuisheng, *Loss properties due to Rayleigh scattering in different types of fiber*, [Optics Express](#) **11**, 39 (2003).
- [72] F. Yu, P. Song, D. Wu, T. Birks, D. Bird, and J. Knight, *Attenuation limit of silica-based hollow-core fiber at mid-IR wavelengths*, [APL Photonics](#) **4**, 080803 (2019).
- [73] M.-J. Li and T. Hayashi, *Advances in low-loss, large-area, and multicore fibers*, in [Optical Fiber Telecommunications VII](#), pp. 3–50 (Elsevier) (2020).
- [74] J. Khurgin and R. Tucker (editors), *Slow light: Science and applications* (CRC Press, Boca Raton) (2009).
- [75] R. M. Camacho, M. V. Pack, J. C. Howell, A. Schweinsberg, and R. W. Boyd, *Wide-bandwidth, tunable, multiple-pulse-width optical delays using slow light in cesium vapor*, [Physical Review Letters](#) **98**, 153601 (2007).
- [76] N. A. Whitaker, M. C. Gabriel, H. Avramopoulos, and A. Huang, *All-optical, all-fiber circulating shift register with an inverter*, [Optics Letters](#) **16**, 1999 (1991).
- [77] H. Avramopoulos and N. A. Whitaker, *Addressable fiber-loop memory*, [Optics Letters](#) **18**, 22 (1993).
- [78] 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).
- [79] X.-L. Pang, A.-L. Yang, J.-P. Dou, H. Li, C.-N. Zhang, E. Poem, D. J. Saunders, H. Tang, J. Nunn, I. A. Walmsley, and X.-M. Jin, *A hybrid quantum memory-enabled network at room temperature*, [Science Advances](#) **6** (2020).
- [80] T. B. Pittman and J. D. Franson, *Cyclical quantum memory for photonic qubits*, [Physical Review A](#) **66**, 062302 (2002).
- [81] P. M. Leung and T. C. Ralph, *Quantum memory scheme based on optical fibers and cavities*, [Physical Review A](#) **74**, 022311 (2006).
- [82] V. Scarani, S. Iblisdir, N. Gisin, and A. Acín, *Quantum cloning*, [Reviews of Modern Physics](#) **77**, 1225 (2005).

- [83] V. Bužek and M. Hillery, *Quantum copying: Beyond the no-cloning theorem*, [Physical Review A](#) **54**, 1844 (1996).
- [84] R. F. Werner, *Optimal cloning of pure states*, [Physical Review A](#) **58**, 1827 (1998).
- [85] N. Gisin, *Quantum cloning without signaling*, [Physics Letters A](#) **242**, 1 (1998).
- [86] N. Gisin and S. Massar, *Optimal quantum cloning machines*, [Physical Review Letters](#) **79**, 2153 (1997).
- [87] F. D. Martini, V. Bužek, F. Sciarrino, and C. Sias, *Experimental realization of the quantum universal NOT gate*, [Nature](#) **419**, 815 (2002).
- [88] M. Ricci, F. Sciarrino, C. Sias, and F. D. Martini, *Teleportation scheme implementing the universal optimal quantum cloning machine and the universal NOT gate*, [Physical Review Letters](#) **92**, 047901 (2004).
- [89] W. T. M. Irvine, A. L. Linares, M. J. A. de Dood, and D. Bouwmeester, *Optimal quantum cloning on a beam splitter*, [Physical Review Letters](#) **92**, 047902 (2004).
- [90] N. J. Cerf, A. Ipe, and X. Rottenberg, *Cloning of continuous quantum variables*, [Physical Review Letters](#) **85**, 1754 (2000).
- [91] F. Grosshans and P. Grangier, *Quantum cloning and teleportation criteria for continuous quantum variables*, [Physical Review A](#) **64**, 010301 (2001).
- [92] P. Sekatski, B. Sanguinetti, E. Pomarico, N. Gisin, and C. Simon, *Cloning entangled photons to scales one can see*, [Physical Review A](#) **82**, 053814 (2010).
- [93] G. M. D'Ariano, F. D. Martini, and M. F. Sacchi, *Continuous variable cloning via network of parametric gates*, [Physical Review Letters](#) **86**, 914 (2001).
- [94] M. K. Tey, Z. Chen, S. A. Aljunid, B. Chng, F. Huber, G. Maslennikov, and C. Kurtsiefer, *Strong interaction between light and a single trapped atom without the need for a cavity*, [Nature Physics](#) **4**, 924 (2008).
- [95] S. J. van Enk and H. J. Kimble, *Strongly focused light beams interacting with single atoms in free space*, [Physical Review A](#) **63**, 023809 (2001).
- [96] M. K. Tey, G. Maslennikov, T. C. H. Liew, S. A. Aljunid, F. Huber, B. Chng, Z. Chen, V. Scarani, and C. Kurtsiefer, *Interfacing light and single atoms with a lens*, [New Journal of Physics](#) **11**, 043011 (2009).
- [97] M. Sondermann, R. Maiwald, H. Konermann, N. Lindlein, U. Peschel, and G. Leuchs, *Design of a mode converter for efficient light-atom coupling in free space*, [Applied Physics B](#) **89**, 489 (2007).
- [98] A. H. Kiilerich and K. Mølmer, *Quantum interactions with pulses of radiation*, [Physical Review A](#) **102**, 023717 (2020).

- [99] J. M. Raimond, M. Brune, and S. Haroche, *Manipulating quantum entanglement with atoms and photons in a cavity*, [Reviews of Modern Physics](#) **73**, 565 (2001).
- [100] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi, *Quantum state transfer and entanglement distribution among distant nodes in a quantum network*, [Physical Review Letters](#) **78**, 3221 (1997).
- [101] A. D. Boozer, A. Boca, R. Miller, T. E. Northup, and H. J. Kimble, *Reversible state transfer between light and a single trapped atom*, [Physical Review Letters](#) **98**, 193601 (2007).
- [102] 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).
- [103] 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).
- [104] P. Törmä, *Transitions in quantum networks*, [Physical Review Letters](#) **81**, 2185 (1998).
- [105] A. Acín, J. I. Cirac, and M. Lewenstein, *Entanglement percolation in quantum networks*, [Nature Physics](#) **3**, 256 (2007).
- [106] M. Afzelius and C. Simon, *Impedance-matched cavity quantum memory*, [Physical Review A](#) **82**, 022310 (2010).
- [107] S. A. Moiseev, S. N. Andrianov, and F. F. Gubaidullin, *Efficient multimode quantum memory based on photon echo in an optimal QED cavity*, [Physical Review A](#) **82**, 022311 (2010).
- [108] M. Sabooni, Q. Li, S. Kröll, and L. Rippe, *Efficient quantum memory using a weakly absorbing sample*, [Physical Review Letters](#) **110**, 133604 (2013).
- [109] P. Jobez, I. Usmani, N. Timoney, C. Laplane, N. Gisin, and M. Afzelius, *Cavity-enhanced storage in an optical spin-wave memory*, [New Journal of Physics](#) **16**, 083005 (2014).
- [110] J. Nunn, J. H. D. Munns, S. Thomas, K. T. Kaczmarek, C. Qiu, A. Feizpour, E. Poem, B. Brecht, D. J. Saunders, P. M. Ledingham, D. V. Reddy, M. G. Raymer, and I. A. Walmsley, *Theory of noise suppression in  $\Lambda$ -type quantum memories by means of a cavity*, [Physical Review A](#) **96**, 012338 (2017).
- [111] A. Wicht, K. Danzmann, M. Fleischhauer, M. Scully, G. Müller, and R.-H. Rinkleff, *White-light cavities, atomic phase coherence, and gravitational wave detectors*, [Optics Communications](#) **134**, 431 (1997).
- [112] R. Fleischhaker and J. Evers, *Four-wave mixing enhanced white-light cavity*, [Physical Review A](#) **78**, 051802 (2008).

- [113] E. S. Moiseev, A. Tashchilina, S. A. Moiseev, and B. C. Sanders, *Broadband quantum memory in a cavity via zero spectral dispersion*, [New Journal of Physics](#) **23**, 063071 (2021).
- [114] 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).
- [115] D. A. Steck, *Rubidium 87 D line data*, available at <https://steck.us/alkalidata> (revision 2.2.2, 2021).
- [116] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in  $\Lambda$ -type optically dense atomic media. I. cavity model*, [Physical Review A](#) **76**, 033804 (2007).
- [117] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in  $\Lambda$ -type optically dense atomic media. II. free-space model*, [Physical Review A](#) **76**, 033805 (2007).
- [118] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, *Photon storage in  $\Lambda$ -type optically dense atomic media. III. effects of inhomogeneous broadening*, [Phys. Rev. A](#) **76**, 033806 (2007).
- [119] A. V. Gorshkov, T. Calarco, M. D. Lukin, and A. S. Sørensen, *Photon storage in  $\Lambda$ -type optically dense atomic media. IV. optimal control using gradient ascent*, [Phys. Rev. A](#) **77**, 043806 (2008).
- [120] M. T. Rakher, R. J. Warburton, and P. Treutlein, *Prospects for storage and retrieval of a quantum-dot single photon in an ultracold  $^{87}\text{Rb}$  ensemble*, [Physical Review A](#) **88**, 053834 (2013).
- [121] A. Dantan, A. Bramati, and M. Pinard, *Atomic quantum memory: Cavity versus single-pass schemes*, [Physical Review A](#) **71**, 043801 (2005).
- [122] E. L. Hahn, *Spin echoes*, [Physical Review](#) **80**, 580 (1950).
- [123] N. A. Kurnit, I. D. Abella, and S. R. Hartmann, *Observation of a photon echo*, [Physical Review Letters](#) **13**, 567 (1964).
- [124] M. Mitsunaga, *Time-domain optical data storage by photon echo*, [Optical and Quantum Electronics](#) **24**, 1137 (1992).
- [125] W. Tittel, M. Afzelius, T. Chanelière, R. L. Cone, S. Kröll, S. A. Moiseev, and M. Sellars, *Photon-echo quantum memory in solid state systems*, [Laser & Photonics Reviews](#) **4**, 244 (2010).
- [126] N. W. Carlson, W. R. Babbitt, T. W. Mossberg, L. J. Rothberg, and A. G. Yodh, *Storage and time reversal of light pulses using photon echoes*, [Optics Letters](#) **8**, 483 (1983).

- [127] T. W. Mossberg, *Time-domain frequency-selective optical data storage*, [Optics Letters](#) **7**, 77 (1982).
- [128] J. Ruggiero, J.-L. L. Gouët, C. Simon, and T. Chanelière, *Why the two-pulse photon echo is not a good quantum memory protocol*, [Physical Review A](#) **79**, 053851 (2009).
- [129] S. A. Moiseev and S. Kröll, *Complete reconstruction of the quantum state of a single-photon wave packet absorbed by a Doppler-broadened transition*, [Physical Review Letters](#) **87**, 173601 (2001).
- [130] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin, *Multimode quantum memory based on atomic frequency combs*, [Physical Review A](#) **79**, 052329 (2009).
- [131] 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).
- [132] 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 (2006).
- [133] 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).
- [134] M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler, *Unconditional room-temperature quantum memory*, [Nature Physics](#) **7**, 794 (2011).
- [135] B. M. Sparkes, J. Bernu, M. Hosseini, J. Geng, Q. Glorieux, P. A. Altin, P. K. Lam, N. P. Robins, and B. C. Buchler, *Gradient echo memory in an ultra-high optical depth cold atomic ensemble*, [New Journal of Physics](#) **15**, 085027 (2013).
- [136] C. W. Thiel, T. Böttger, and R. L. Cone, *Rare-earth-doped materials for applications in quantum information storage and signal processing*, [Journal of Luminescence](#) **131**, 353 (2011).
- [137] E. Fraval, M. J. Sellars, and J. J. Longdell, *Dynamic decoherence control of a solid-state nuclear-quadrupole qubit*, [Physical Review Letters](#) **95**, 030506 (2005).
- [138] E. Fraval, M. J. Sellars, and J. J. Longdell, *Method of extending hyperfine coherence times in  $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$* , [Physical Review Letters](#) **92**, 077601 (2004).
- [139] M. Zhong, M. P. Hedges, R. L. Ahlefeldt, J. G. Bartholomew, S. E. Beavan, S. M. Wittig, J. J. Longdell, and M. J. Sellars, *Optically addressable nuclear spins in a solid with a six-hour coherence time*, [Nature](#) **517**, 177 (2015).
- [140] G. Hétet, J. J. Longdell, A. L. Alexander, P. K. Lam, and M. J. Sellars, *Electro-optic quantum memory for light using two-level atoms*, [Physical Review Letters](#) **100**, 023601 (2008).

- [141] M. P. Hedges, J. J. Longdell, Y. Li, and M. J. Sellars, *Efficient quantum memory for light*, [Nature](#) **465**, 1052 (2010).
- [142] S. P. Horvath, M. K. Alqedra, A. Kinos, A. Walther, J. M. Dahlström, S. Kröll, and L. Rippe, *Noise-free on-demand atomic frequency comb quantum memory*, [Physical Review Research](#) **3**, 023099 (2021).
- [143] M. Nilsson, L. Rippe, S. Kröll, R. Klieber, and D. Suter, *Hole-burning techniques for isolation and study of individual hyperfine transitions in inhomogeneously broadened solids demonstrated in  $Pr^{3+}:Y_2SiO_5$* , [Physical Review B](#) **70**, 214116 (2004).
- [144] N. Sangouard, C. Simon, M. Afzelius, and N. Gisin, *Analysis of a quantum memory for photons based on controlled reversible inhomogeneous broadening*, [Physical Review A](#) **75**, 032327 (2007).
- [145] S. A. Moiseev and N. M. Arslanov, *Efficiency and fidelity of photon-echo quantum memory in an atomic system with longitudinal inhomogeneous broadening*, [Physical Review A](#) **78**, 023803 (2008).
- [146] J. Nunn, K. Reim, K. C. Lee, V. O. Lorenz, B. J. Sussman, I. A. Walmsley, and D. Jaksch, *Multimode memories in atomic ensembles*, [Physical Review Letters](#) **101**, 260502 (2008).
- [147] 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).
- [148] 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).
- [149] T. Chanelière, J. Ruggiero, M. Bonarota, M. Afzelius, and J.-L. L. Gouët, *Efficient light storage in a crystal using an atomic frequency comb*, [New Journal of Physics](#) **12**, 023025 (2010).
- [150] M. Rančić, M. P. Hedges, R. L. Ahlefeldt, and M. J. Sellars, *Coherence time of over a second in a telecom-compatible quantum memory storage material*, [Nature Physics](#) **14**, 50 (2017).
- [151] D. Main, T. M. Hird, S. Gao, I. A. Walmsley, and P. M. Ledingham, *Room temperature atomic frequency comb storage for light*, [Optics Letters](#) **46**, 2960 (2021).
- [152] J. L. Rubio, D. Viscor, J. Mompert, and V. Ahufinger, *Atomic-frequency-comb quantum memory via piecewise adiabatic passage*, [Physical Review A](#) **98**, 043834 (2018).
- [153] S. E. Harris, *Electromagnetically induced transparency*, [Physics Today](#) **50**, 36 (1997).

- [154] S. E. Harris, J. E. Field, and A. Imamoglu, *Nonlinear optical processes using electromagnetically induced transparency*, [Physical Review Letters](#) **64**, 1107 (1990).
- [155] K.-J. Boller, A. Imamoglu, and S. E. Harris, *Observation of electromagnetically induced transparency*, [Physical Review Letters](#) **66**, 2593 (1991).
- [156] M. Fleischhauer, A. Imamoglu, and J. P. Marangos, *Electromagnetically induced transparency: Optics in coherent media*, [Reviews of Modern Physics](#) **77**, 633 (2005).
- [157] S. E. Harris, J. E. Field, and A. Kasapi, *Dispersive properties of electromagnetically induced transparency*, [Physical Review A](#) **46**, R29 (1992).
- [158] 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).
- [159] R. W. Boyd and D. J. Gauthier, *Progress in optics*, in E. Wolf (editor), *Progress in Optics*, volume 43, chapter “Slow” and “fast” light, pp. 497–530 (Elsevier Science B.V.), first edition (2002).
- [160] C. N. Cohen-Tannoudji, *The Autler-Townes effect revisited*, in *Amazing light*, pp. 109–123 (Springer New York) (1996).
- [161] T. F. Krauss, *Slow light in photonic crystal waveguides*, [Journal of Physics D: Applied Physics](#) **40**, 2666 (2007).
- [162] 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).
- [163] M. Fleischhauer and M. D. Lukin, *Quantum memory for photons: Dark-state polaritons*, [Physical Review A](#) **65**, 022314 (2002).
- [164] J. Oreg, F. T. Hioe, and J. H. Eberly, *Adiabatic following in multilevel systems*, [Physical Review A](#) **29**, 690 (1984).
- [165] M. D. Lukin, *Colloquium: Trapping and manipulating photon states in atomic ensembles*, [Reviews of Modern Physics](#) **75**, 457 (2003).
- [166] 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).
- [167] M. D. Lukin, M. Fleischhauer, A. S. Zibrov, H. G. Robinson, V. L. Velichansky, L. Hollberg, and M. O. Scully, *Spectroscopy in dense coherent media: Line narrowing and interference effects*, [Physical Review Letters](#) **79**, 2959 (1997).



- [168] N. B. Phillips, A. V. Gorshkov, and I. Novikova, *Optimal light storage in atomic vapor*, [Physical Review A](#) **78**, 023801 (2008).
- [169] 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).
- [170] C. Kupchak, T. Mittiga, B. Jordaan, M. Namazi, C. Nölleke, and E. Figueroa, *Room-temperature single-photon level memory for polarization states*, [Scientific Reports](#) **5** (2015).
- [171] M. Dąbrowski, R. Chrapkiewicz, and W. Wasilewski, *Magnetically tuned, robust and efficient filtering system for spatially multimode quantum memory in warm atomic vapors*, [Journal of Modern Optics](#) **63**, 2029 (2015).
- [172] 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).
- [173] 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).
- [174] 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).
- [175] 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).
- [176] 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).
- [177] 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).
- [178] D. G. England, K. A. Fisher, J.-P. W. MacLean, P. J. Bustard, R. Lausten, K. J. Resch, and B. J. Sussman, *Storage and retrieval of THz-bandwidth single photons using a room-temperature diamond quantum memory*, [Physical Review Letters](#) **114**, 053602 (2015).
- [179] D. L. Rousseau, G. D. Patterson, and P. F. Williams, *Resonance Raman Scattering and Collision-Induced Redistribution Scattering in I<sub>2</sub>*, [Physical Review Letters](#) **34**, 1306 (1975).

- [180] M. T. L. Hsu, G. Hétet, O. Glöckl, J. J. Longdell, B. C. Buchler, H.-A. Bachor, and P. K. Lam, *Quantum study of information delay in electromagnetically induced transparency*, [Physical Review Letters](#) **97**, 183601 (2006).
- [181] S. Manz, T. Fernholz, J. Schmiedmayer, and J.-W. Pan, *Collisional decoherence during writing and reading quantum states*, [Physical Review A](#) **75**, 040101 (2007).
- [182] 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).
- [183] G. S. Agarwal, T. N. Dey, and D. J. Gauthier, *Competition between electromagnetically induced transparency and Raman processes*, [Physical Review A](#) **74**, 043805 (2006).
- [184] 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).
- [185] 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).
- [186] 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).
- [187] 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).
- [188] G. Heinze, C. Hubrich, and T. Halfmann, *Stopped light and image storage by electromagnetically induced transparency up to the regime of one minute*, [Physical Review Letters](#) **111**, 033601 (2013).
- [189] Y. O. Dudin, L. Li, and A. Kuzmich, *Light storage on the time scale of a minute*, [Physical Review A](#) **87**, 031801 (2013).
- [190] H. Chi, W. Quan, J. Zhang, L. Zhao, and J. Fang, *Advances in anti-relaxation coatings of alkali-metal vapor cells*, [Applied Surface Science](#) **501**, 143897 (2020).
- [191] M. V. Balabas, K. Jensen, W. Wasilewski, H. Krauter, L. S. Madsen, J. H. Müller, T. Fernholz, and E. S. Polzik, *High quality anti-relaxation coating material for alkali atom vapor cells*, [Optics Express](#) **18**, 5825 (2010).

- [192] M. V. Balabas, T. Karaulanov, M. P. Ledbetter, and D. Budker, *Polarized alkali-metal vapor with minute-long transverse spin-relaxation time*, [Physical Review Letters](#) **105**, 070801 (2010).
- [193] E. P. Corsini, T. Karaulanov, M. Balabas, and D. Budker, *Hyperfine frequency shift and Zeeman relaxation in alkali-metal-vapor cells with antirelaxation alkene coating*, [Physical Review A](#) **87**, 022901 (2013).
- [194] W. Happer, *Optical pumping*, [Reviews of Modern Physics](#) **44**, 169 (1972).
- [195] J. Vanier and C. Audoin, *The quantum physics of atomic frequency standards* (IOP Publishing Ltd) (1989).
- [196] D. Budker and M. Romalis, *Optical magnetometry*, [Nature Physics](#) **3**, 227 (2007).
- [197] O. Katz and O. Firstenberg, *Light storage for one second in room-temperature alkali vapor*, [Nature Communications](#) **9** (2018).
- [198] J. Nunn, *Quantum memory in atomic ensembles*, [Ph.D. thesis](#), St. John's College, University of Oxford (2008).
- [199] A. E. Kozhokin, K. Mølmer, and E. Polzik, *Quantum memory for light*, [Physical Review A](#) **62**, 033809 (2000).
- [200] J. Nunn, I. A. Walmsley, M. G. Raymer, K. Surmacz, F. C. Waldermann, Z. Wang, and D. Jaksch, *Mapping broadband single-photon wave packets into an atomic memory*, [Physical Review A](#) **75**, 011401 (2007).
- [201] 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).
- [202] G. Romanov, C. O'Brien, and I. Novikova, *Suppression of the four-wave mixing amplification via Raman absorption*, [Journal of Modern Optics](#) **63**, 2048 (2016).
- [203] S. E. Thomas, T. M. Hird, J. H. D. Munns, B. Brecht, D. J. Saunders, J. Nunn, I. A. Walmsley, and P. M. Ledingham, *Raman quantum memory with built-in suppression of four-wave-mixing noise*, [Physical Review A](#) **100**, 033801 (2019).
- [204] P. J. Bustard, D. G. England, K. Heshami, C. Kupchak, and B. J. Sussman, *Reducing noise in a Raman quantum memory*, [Optics Letters](#) **41**, 5055 (2016).
- [205] K. A. G. Fisher, D. G. England, J.-P. W. MacLean, P. J. Bustard, K. J. Resch, and B. J. Sussman, *Frequency and bandwidth conversion of single photons in a room-temperature diamond quantum memory*, [Nature Communications](#) **7** (2016).
- [206] 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).

- [207] 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).
- [208] 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).
- [209] 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).
- [210] W. S. Martins, D. M. Conrado, V. Ádony, S. Barreiro, M. Oriá, M. Chevrollier, and R. A. de Oliveira, *Efficient atomic memory using electromagnetically induced absorption*, [Physics Open](#) **9**, 100081 (2021).
- [211] K. Pandey, *Role of different types of subsystems in a doubly driven  $\Lambda$  system in  $^{87}\text{Rb}$* , [Physical Review A](#) **87**, 043838 (2013).
- [212] J. M. Taylor, C. M. Marcus, and M. D. Lukin, *Long-lived memory for mesoscopic quantum bits*, [Physical Review Letters](#) **90**, 206803 (2003).
- [213] T. R. Gentile, P. J. Nacher, B. Saam, and T. G. Walker, *Optically polarized  $^3\text{He}$* , [Reviews of Modern Physics](#) **89**, 045004 (2017).
- [214] A. Dantan, G. Reinaudi, A. Sinatra, F. Laloë, E. Giacobino, and M. Pinard, *Long-lived quantum memory with nuclear atomic spins*, [Physical Review Letters](#) **95**, 123002 (2005).
- [215] A. Serafin, M. Fadel, P. Treutlein, and A. Sinatra, *Nuclear spin squeezing in helium-3 by continuous quantum nondemolition measurement*, [Physical Review Letters](#) **127**, 013601 (2021).
- [216] S. L. Braunstein and P. van Loock, *Quantum information with continuous variables*, [Reviews of Modern Physics](#) **77**, 513 (2005).
- [217] K. Hammerer, A. S. Sørensen, and E. S. Polzik, *Quantum interface between light and atomic ensembles*, [Reviews of Modern Physics](#) **82**, 1041 (2010).
- [218] F. Flamini, N. Spagnolo, and F. Sciarrino, *Photonic quantum information processing: A review*, [Reports on Progress in Physics](#) **82**, 016001 (2019).
- [219] B. Julsgaard, J. Sherson, J. I. Cirac, J. Fiurášek, and E. S. Polzik, *Experimental demonstration of quantum memory for light*, [Nature](#) **432**, 482 (2004).
- [220] O. Katz, R. Shaham, and O. Firstenberg, *Quantum interface for noble-gas spins based on spin-exchange collisions*, [PRX Quantum](#) **3**, 010305 (2022).

- [221] R. Shaham, O. Katz, and O. Firstenberg, *Strong coupling of alkali-metal spins to noble-gas spins with an hour-long coherence time*, [Nature Physics](#) **18**, 506 (2022).
- [222] O. Katz, R. Shaham, E. Reches, A. V. Gorshkov, and O. Firstenberg, *Optical quantum memory for noble-gas spins based on spin-exchange collisions*, [Physical Review A](#) **105**, 042606 (2022).
- [223] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Photons and atoms: Introduction to quantum electrodynamics* (Wiley-Interscience, New York) (1997).
- [224] U. M. Titulaer and R. J. Glauber, *Density operators for coherent fields*, [Physical Review](#) **145**, 1041 (1966).
- [225] B. J. Smith and M. G. Raymer, *Photon wave functions, wave-packet quantization of light, and coherence theory*, [New Journal of Physics](#) **9**, 414 (2007).
- [226] L. Mandel and E. Wolf, *Optical coherence and quantum optics*, (Cambridge University Press) (1995).
- [227] 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).
- [228] B. Lounis and M. Orrit, *Single-photon sources*, [Reports on Progress in Physics](#) **68**, 1129 (2005).
- [229] I. Aharonovich, D. Englund, and M. Toth, *Solid-state single-photon emitters*, [Nature Photonics](#) **10**, 631 (2016).
- [230] R. Boyd, *Nonlinear optics (third edition)* (Academic Press) (2008).
- [231] P. E. Powers and J. W. Haus, *Fundamentals of nonlinear optics* (CRC Press) (2019).
- [232] M. C. Tichy, M. Tiersch, F. de Melo, F. Mintert, and A. Buchleitner, *Zero-transmission law for multiport beam splitters*, [Physical Review Letters](#) **104**, 220405 (2010).
- [233] A. Crespi, *Suppression laws for multiparticle interference in Sylvester interferometers*, [Physical Review A](#) **91**, 013811 (2015).
- [234] N. Viggianiello, F. Flamini, M. Bentivegna, N. Spagnolo, A. Crespi, D. J. Brod, E. F. Galvão, R. Osellame, and F. Sciarrino, *Optimal photonic indistinguishability tests in multimode networks*, [Science Bulletin](#) **63**, 1470 (2018).
- [235] T. Wasak, A. Smerzi, and J. Chwedeńczuk, *Role of particle entanglement in the violation of bell inequalities*, [Scientific Reports](#) **8** (2018).

- [236] 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).
- [237] R. J. Glauber, *The quantum theory of optical coherence*, [Physical Review](#) **130**, 2529 (1963).
- [238] R. J. Glauber, *Coherent and incoherent states of the radiation field*, [Physical Review](#) **131**, 2766 (1963).
- [239] P. Senellart, G. Solomon, and A. White, *High-performance semiconductor quantum-dot single-photon sources*, [Nature Nanotechnology](#) **12**, 1026 (2017).
- [240] 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).
- [241] 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).
- [242] 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).
- [243] A. Christ and C. Silberhorn, *Limits on the deterministic creation of pure single-photon states using parametric down-conversion*, [Physical Review A](#) **85**, 023829 (2012).
- [244] E. Meyer-Scott, C. Silberhorn, and A. Migdall, *Single-photon sources: Approaching the ideal through multiplexing*, [Review of Scientific Instruments](#) **91**, 041101 (2020).
- [245] J. F. Clauser, *Experimental distinction between the quantum and classical field-theoretic predictions for the photoelectric effect*, [Physical Review D](#) **9**, 853 (1974).
- [246] J. F. Clauser and A. Shimony, *Bell's theorem. experimental tests and implications*, [Reports on Progress in Physics](#) **41**, 1881 (1978).
- [247] A. Aspect, P. Grangier, and G. Roger, *Experimental tests of realistic local theories via Bell's theorem*, [Physical Review Letters](#) **47**, 460 (1981).
- [248] A. Aspect, J. Dalibard, and G. Roger, *Experimental test of Bell's inequalities using time-varying analyzers*, [Physical Review Letters](#) **49**, 1804 (1982).
- [249] J. McKeever, *Deterministic generation of single photons from one atom trapped in a cavity*, [Science](#) **303**, 1992 (2004).
- [250] D. B. Higginbottom, L. Slodička, G. Araneda, L. Lachman, R. Filip, M. Hennrich, and R. Blatt, *Pure single photons from a trapped atom source*, [New Journal of Physics](#) **18**, 093038 (2016).

- [251] A. MacRae, T. Brannan, R. Achal, and A. I. Lvovsky, *Tomography of a high-purity narrowband photon from a transient atomic collective excitation*, [Physical Review Letters](#) **109**, 033601 (2012).
- [252] 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).
- [253] M. Afzelius, N. Gisin, and H. de Riedmatten, *Quantum memory for photons*, [Physics Today](#) **68**, 42 (2015).
- [254] 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).
- [255] 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).
- [256] N. Maring, P. Farrera, K. Kutluer, M. Mazzer, G. Heinze, and H. de Riedmatten, *Photonic quantum state transfer between a cold atomic gas and a crystal*, [Nature](#) **551**, 485 (2017).
- [257] D. Leonard, K. Pond, and P. M. Petroff, *Critical layer thickness for self-assembled InAs islands on GaAs*, [Physical Review B](#) **50**, 11687 (1994).
- [258] R. J. Warburton, *Single spins in self-assembled quantum dots*, [Nature Materials](#) **12**, 483 (2013).
- [259] Y. H. Huo, A. Rastelli, and O. G. Schmidt, *Ultra-small excitonic fine structure splitting in highly symmetric quantum dots on GaAs (001) substrate*, [Applied Physics Letters](#) **102**, 152105 (2013).
- [260] N. Akopian, L. Wang, A. Rastelli, O. G. Schmidt, and V. Zwiller, *Hybrid semiconductor-atomic interface: slowing down single photons from a quantum dot*, [Nature Photonics](#) **5**, 230 (2011).
- [261] T. Kroh, J. Wolters, A. Ahlrichs, A. W. Schell, A. Thoma, S. Reitzenstein, J. S. Wildmann, E. Zallo, R. Trotta, A. Rastelli, O. G. Schmidt, and O. Benson, *Slow and fast single photons from a quantum dot interacting with the excited state hyperfine structure of the cesium  $D_1$ -line*, [Scientific Reports](#) **9** (2019).
- [262] H. Vural, S. Seyfferle, I. Gerhardt, M. Jetter, S. L. Portalupi, and P. Michler, *Delaying two-photon Fock states in hot cesium vapor using single photons generated on demand from a semiconductor quantum dot*, [Physical Review B](#) **103**, 195304 (2021).
- [263] 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).
- [264] T. Lettner, K. D. Zeuner, E. Schöll, H. Huang, S. Scharmer, S. F. C. da Silva, S. Gyger, L. Schweickert, A. Rastelli, K. D. Jöns, and V. Zwiller, *GaAs quantum dot in a parabolic microcavity tuned to  $^{87}\text{Rb}$   $D_1$* , [ACS Photonics](#) **7**, 29 (2019).
- [265] A. Högele, S. Seidl, M. Kroner, K. Karrai, R. J. Warburton, B. D. Gerardot, and P. M. Petroff, *Voltage-controlled optics of a quantum dot*, [Physical Review Letters](#) **93**, 217401 (2004).
- [266] A. V. Kuhlmann, J. Houel, A. Ludwig, L. Greuter, D. Reuter, A. D. Wieck, M. Poggio, and R. J. Warburton, *Charge noise and spin noise in a semiconductor quantum device*, [Nature Physics](#) **9**, 570 (2013).
- [267] S. Kumar, R. Trotta, E. Zallo, J. D. Plumhof, P. Atkinson, A. Rastelli, and O. G. Schmidt, *Strain-induced tuning of the emission wavelength of high quality GaAs/AlGaAs quantum dots in the spectral range of the  $^{87}\text{Rb}$   $D_2$  lines*, [Applied Physics Letters](#) **99**, 161118 (2011).
- [268] 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).
- [269] 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).
- [270] 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** (2020).
- [271] C. Matthiesen, M. Geller, C. H. H. Schulte, C. L. Gall, J. Hansom, Z. Li, M. Hugues, E. Clarke, and M. Atatüre, *Phase-locked indistinguishable photons with synthesized waveforms from a solid-state source*, [Nature Communications](#) **4** (2013).
- [272] 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).
- [273] L. Schweickert, K. D. Jöns, K. D. Zeuner, S. F. C. da Silva, H. Huang, T. Lettner, M. Reindl, J. Zichi, R. Trotta, A. Rastelli, and V. Zwiller, *On-demand generation of background-free single photons from a solid-state source*, [Applied Physics Letters](#) **112**, 093106 (2018).



- [274] 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).
- [275] H. Wang, Y. He, Y.-H. Li, Z.-E. Su, B. Li, H.-L. Huang, X. Ding, M.-C. Chen, C. Liu, J. Qin, J.-P. Li, Y.-M. He, C. Schneider, M. Kamp, C.-Z. Peng, S. Höfling, C.-Y. Lu, and J.-W. Pan, *High-efficiency multiphoton boson sampling*, [Nature Photonics](#) **11**, 361 (2017).
- [276] R. B. Patel, A. J. Bennett, I. Farrer, C. A. Nicoll, D. A. Ritchie, and A. J. Shields, *Two-photon interference of the emission from electrically tunable remote quantum dots*, [Nature Photonics](#) **4**, 632 (2010).
- [277] V. Giesz, S. L. Portalupi, T. Grange, C. Antón, L. D. Santis, J. Demory, N. Somaschi, I. Sagnes, A. Lemaître, L. Lanco, A. Auffèves, and P. Senellart, *Cavity-enhanced two-photon interference using remote quantum dot sources*, [Physical Review B](#) **92**, 161302 (2015).
- [278] M. Reindl, K. D. Jöns, D. Huber, C. Schimpf, Y. Huo, V. Zwiller, A. Rastelli, and R. Trotta, *Phonon-assisted two-photon interference from remote quantum emitters*, [Nano Letters](#) **17**, 4090 (2017).
- [279] E. Schöll, L. Hanschke, L. Schweickert, K. D. Zeuner, M. Reindl, S. F. C. da Silva, T. Lettner, R. Trotta, J. J. Finley, K. Müller, A. Rastelli, V. Zwiller, and K. D. Jöns, *Resonance fluorescence of GaAs quantum dots with near-unity photon indistinguishability*, [Nano Letters](#) **19**, 2404 (2019).
- [280] 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](#) (2022).
- [281] 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).
- [282] S. Emanuelli and A. Arie, *Temperature-dependent dispersion equations for  $KTiOPO_4$  and  $KTiOAsO_4$* , [Applied Optics](#) **42**, 6661 (2003).
- [283] 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).
- [284] Y. J. Lu and Z. Y. Ou, *Optical parametric oscillator far below threshold: Experiment versus theory*, [Physical Review A](#) **62**, 033804 (2000).

- [285] Y. Jeronimo-Moreno, S. Rodriguez-Benavides, and A. B. U'Ren, *Theory of cavity-enhanced spontaneous parametric downconversion*, [Laser Physics](#) **20**, 1221 (2010).
- [286] M. Rambach, A. Nikolova, T. J. Weinhold, and A. G. White, *Sub-megahertz linewidth single photon source*, [APL Photonics](#) **1**, 096101 (2016).
- [287] 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).
- [288] 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).
- [289] N. Gisin and R. Thew, *Quantum communication*, [Nature Photonics](#) **1**, 165 (2007).
- [290] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, in *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*, pp. 175–179 (1984).
- [291] 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).
- [292] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, *Fundamental limits of repeaterless quantum communications*, [Nature Communications](#) **8** (2017).
- [293] R. Valivarthi, M. G. Puigibert, Q. Zhou, G. H. Aguilar, V. B. Verma, F. Marsili, M. D. Shaw, S. W. Nam, D. Oblak, and W. Tittel, *Quantum teleportation across a metropolitan fibre network*, [Nature Photonics](#) **10**, 676 (2016).
- [294] 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).
- [295] 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).
- [296] Z.-D. Li, R. Zhang, X.-F. Yin, L.-Z. Liu, Y. Hu, Y.-Q. Fang, Y.-Y. Fei, X. Jiang, J. Zhang, L. Li, N.-L. Liu, F. Xu, Y.-A. Chen, and J.-W. Pan, *Experimental quantum repeater without quantum memory*, [Nature Photonics](#) **13**, 644 (2019).
- [297] 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).

- [298] F. Monteiro, V. C. Vivoli, T. Guerreiro, A. Martin, J.-D. Bancal, H. Zbinden, R. Thew, and N. Sangouard, *Revealing genuine optical-path entanglement*, [Physical Review Letters](#) **114**, 170504 (2015).
- [299] V. C. Vivoli, P. Sekatski, J.-D. Bancal, C. C. W. Lim, A. Martin, R. T. Thew, H. Zbinden, N. Gisin, and N. Sangouard, *Comparing different approaches for generating random numbers device-independently using a photon pair source*, [New Journal of Physics](#) **17**, 023023 (2015).
- [300] 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** (2017).
- [301] M. Lipka, M. Mazelanik, and M. Parniak, *Entanglement distribution with wavevector-multiplexed quantum memory*, [New Journal of Physics](#) **23**, 053012 (2021).
- [302] M. Parniak, M. Dąbrowski, M. Mazelanik, A. Leszczyński, M. Lipka, and W. Wasilewski, *Wavevector multiplexed atomic quantum memory via spatially-resolved single-photon detection*, [Nature Communications](#) **8** (2017).
- [303] N. Sinclair, E. Saglamyurek, H. Mallahzadeh, J. A. Slater, M. George, R. Ricken, M. P. Hedges, D. Oblak, C. Simon, W. Sohler, and W. Tittel, *Spectral multiplexing for scalable quantum photonics using an atomic frequency comb quantum memory and feed-forward control*, [Physical Review Letters](#) **113**, 053603 (2014).
- [304] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, *Mixed-state entanglement and quantum error correction*, [Physical Review A](#) **54**, 3824 (1996).
- [305] J.-W. Pan, C. Simon, Č. Brukner, and A. Zeilinger, *Entanglement purification for quantum communication*, [Nature](#) **410**, 1067 (2001).
- [306] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, *Entanglement distillation between solid-state quantum network nodes*, [Science](#) **356**, 928 (2017).
- [307] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, *Quantum cryptography*, [Reviews of Modern Physics](#) **74**, 145 (2002).
- [308] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, *The security of practical quantum key distribution*, [Reviews of Modern Physics](#) **81**, 1301 (2009).
- [309] 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).
- [310] M. Herrero-Collantes and J. C. Garcia-Escartin, *Quantum random number generators*, [Reviews of Modern Physics](#) **89**, 015004 (2017).

- [311] C. Portmann and R. Renner, *Security in quantum cryptography*, [Reviews of Modern Physics](#) **94**, 025008 (2022).
- [312] C. H. Bennett, G. Brassard, C. Crepeau, and U. M. Maurer, *Generalized privacy amplification*, [IEEE Transactions on Information Theory](#) **41**, 1915 (1995).
- [313] I. Gerhardt, Q. Liu, A. Lamas-Linares, J. Skaar, C. Kurtsiefer, and V. Makarov, *Full-field implementation of a perfect eavesdropper on a quantum cryptography system*, [Nature Communications](#) **2** (2011).
- [314] V. Makarov, J.-P. Bourgoin, P. Chaiwongkhot, M. Gagné, T. Jennewein, S. Kaiser, R. Kashyap, M. Legré, C. Minshull, and S. Sajeed, *Creation of backdoors in quantum communications via laser damage*, [Physical Review A](#) **94**, 030302 (2016).
- [315] S. Sajeed, P. Chaiwongkhot, A. Huang, H. Qin, V. Egorov, A. Kozubov, A. Gaidash, V. Chistiakov, A. Vasiliev, A. Gleim, and V. Makarov, *An approach for security evaluation and certification of a complete quantum communication system*, [Scientific Reports](#) **11** (2021).
- [316] P. W. Shor, *Algorithms for quantum computation: discrete logarithms and factoring*, in [Proceedings 35th Annual Symposium on Foundations of Computer Science \(IEEE Comput. Soc. Press\)](#) (1994).
- [317] C. Gidney and M. Ekerå, *How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits*, [Quantum](#) **5**, 433 (2021).
- [318] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Surface codes: Towards practical large-scale quantum computation*, [Physical Review A](#) **86**, 032324 (2012).
- [319] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, *Strengths and weaknesses of quantum computing*, [SIAM Journal on Computing](#) **26**, 1510 (1997).
- [320] R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven, *Focus beyond quadratic speedups for error-corrected quantum advantage*, [Physical Review X Quantum](#) **2**, 010103 (2021).
- [321] M. Kaplan, G. Leurent, A. Leverrier, and M. Naya-Plasencia, *Quantum differential and linear cryptanalysis*, [IACR Transactions on Symmetric Cryptology](#) pp. 71–94 (2016).
- [322] D. J. Bernstein, J. Buchmann, and E. Dahmen (editors), *Post-quantum cryptography*, ([Springer Berlin Heidelberg](#)) (2009).
- [323] G. Alagic, J. Alperin-Sheriff, D. C. Apon, D. A. Cooper, Q. H. Dang, J. M. Kelsey, Y.-K. Liu, C. A. Miller, D. Moody, R. C. Peralta, R. A. Perlner, A. Y. Robinson, and D. C. Smith-Tone, *Status report on the second round of the NIST post-quantum cryptography standardization process*, [Technical report, National Institute of Standards and Technology](#) (2020).

- [324] M. Curty and D. J. Santos, *Quantum authentication of classical messages*, [Physical Review A](#) **64**, 062309 (2001).
- [325] R. J. Donaldson, R. J. Collins, K. Kleczkowska, R. Amiri, P. Wallden, V. Dunjko, J. Jeffers, E. Andersson, and G. S. Buller, *Experimental demonstration of kilometer-range quantum digital signatures*, [Physical Review A](#) **93**, 012329 (2016).
- [326] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, *Quantum fingerprinting*, [Physical Review Letters](#) **87**, 167902 (2001).
- [327] X. Zhong, F. Xu, H.-K. Lo, and L. Qian, *Efficient experimental quantum fingerprinting with channel multiplexing and simultaneous detection*, [Nature Communications](#) **12** (2021).
- [328] C. Thalacker, F. Hahn, J. de Jong, A. Pappa, and S. Barz, *Anonymous and secret communication in quantum networks*, [New Journal of Physics](#) **23**, 083026 (2021).
- [329] G. Brassard, A. Broadbent, J. Fitzsimons, S. Gambs, and A. Tapp, *Anonymous quantum communication*, in K. Kurosawa (editor), *Advances in Cryptology – ASIACRYPT 2007*, pp. 460–473 (Springer Berlin Heidelberg, Berlin, Heidelberg) (2007).
- [330] H.-K. Lo and H. F. Chau, *Is quantum bit commitment really possible?*, [Physical Review Letters](#) **78**, 3410 (1997).
- [331] D. Mayers, *Unconditionally secure quantum bit commitment is impossible*, [Physical Review Letters](#) **78**, 3414 (1997).
- [332] H. Buhrman, M. Christandl, P. Hayden, H.-K. Lo, and S. Wehner, *Security of quantum bit string commitment depends on the information measure*, [Physical Review Letters](#) **97**, 250501 (2006).
- [333] G. P. He, *Quantum key distribution based on orthogonal states allows secure quantum bit commitment*, [Journal of Physics A: Mathematical and Theoretical](#) **44**, 445305 (2011).
- [334] M. Heuck, K. Jacobs, and D. R. Englund, *Controlled-phase gate using dynamically coupled cavities and optical nonlinearities*, [Physical Review Letters](#) **124**, 160501 (2020).
- [335] Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble, *Measurement of conditional phase shifts for quantum logic*, [Physical Review Letters](#) **75**, 4710 (1995).
- [336] D. J. Brod and J. Combes, *Passive CPHASE gate via cross-Kerr nonlinearities*, [Physical Review Letters](#) **117**, 080502 (2016).
- [337] S. D. Barrett, P. P. Rohde, and T. M. Stace, *Scalable quantum computing with atomic ensembles*, [New Journal of Physics](#) **12**, 093032 (2010).

- [338] K. M. Beck, M. Hosseini, Y. Duan, and V. Vuletić, *Large conditional single-photon cross-phase modulation*, [Proceedings of the National Academy of Sciences](#) **113**, 9740 (2016).
- [339] K. Koshino, S. Ishizaka, and Y. Nakamura, *Deterministic photon-photon  $\sqrt{SWAP}$  gate using a  $\Lambda$  system*, [Physical Review A](#) **82**, 010301 (2010).
- [340] H. J. Briegel and R. Raussendorf, *Persistent entanglement in arrays of interacting particles*, [Physical Review Letters](#) **86**, 910 (2001).
- [341] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, *Linear optical quantum computing with photonic qubits*, [Reviews of Modern Physics](#) **79**, 135 (2007).
- [342] L.-M. Duan and R. Raussendorf, *Efficient quantum computation with probabilistic quantum gates*, [Physical Review Letters](#) **95**, 080503 (2005).
- [343] J. Borregaard, A. S. Sørensen, J. I. Cirac, and M. D. Lukin, *Efficient quantum computation in a network with probabilistic gates and logical encoding*, [Physical Review A](#) **95**, 042312 (2017).
- [344] C. Simon, *Towards a global quantum network*, [Nature Photonics](#) **11**, 678 (2017).
- [345] K. Makino, Y. Hashimoto, J. ichi Yoshikawa, H. Ohdan, T. Toyama, P. van Loock, and A. Furusawa, *Synchronization of optical photons for quantum information processing*, [Science Advances](#) **2** (2016).
- [346] 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](#) (ACM Press) (2011).
- [347] M. Tillmann, B. Dakić, R. Heilmann, S. Nolte, A. Szameit, and P. Walther, *Experimental boson sampling*, [Nature Photonics](#) **7**, 540 (2013).
- [348] J. Wu, Y. Liu, B. Zhang, X. Jin, Y. Wang, H. Wang, and X. Yang, *A benchmark test of boson sampling on Tianhe-2 supercomputer*, [National Science Review](#) **5**, 715 (2018).
- [349] L. Chakhmakhchyan and N. J. Cerf, *Boson sampling with Gaussian measurements*, [Physical Review A](#) **96**, 032326 (2017).
- [350] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, *Quantum computational advantage using photons*, [Science](#) **370**, 1460 (2020).
- [351] J. F. F. Bulmer, B. A. Bell, R. S. Chadwick, A. E. Jones, D. Moise, A. Rigazzi, J. Thorbecke, U.-U. Haus, T. V. Vaerenbergh, R. B. Patel, I. A. Walmsley, and

- A. Laing, *The boundary for quantum advantage in Gaussian boson sampling*, [Science Advances](#) **8** (2022).
- [352] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien, *Quantum computers*, [Nature](#) **464**, 45 (2010).
- [353] R. Van Meter and S. J. Devitt, *The path to scalable distributed quantum computing*, [Computer](#) **49**, 31 (2016).
- [354] 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*, 16, pp. 1–13 (ACM) (2021).
- [355] L. K. Grover, *Quantum teleportation*, [arXiv:quant-ph/9704012](#) (1997).
- [356] D. Cuomo, M. Caleffi, and A. S. Cacciapuoti, *Towards a distributed quantum computing ecosystem*, [IET Quantum Communication](#) **1**, 3 (2020).
- [357] L. Lamport, R. Shostak, and M. Pease, *The Byzantine generals problem*, [ACM Transactions on Programming Languages and Systems](#) **4**, 382 (1982).
- [358] M. Fitzi, N. Gisin, and U. Maurer, *Quantum solution to the Byzantine agreement problem*, [Physical Review Letters](#) **87**, 217901 (2001).
- [359] M. Fitzi, N. Gisin, U. Maurer, and O. von Rotz, *Unconditional Byzantine agreement and multi-party computation secure against dishonest minorities from scratch*, in L. R. Knudsen (editor), *Advances in Cryptology — EUROCRYPT 2002*, pp. 482–501 (Springer Berlin Heidelberg, Berlin, Heidelberg) (2002).
- [360] M. Ben-Or and A. Hassidim, *Fast quantum Byzantine agreement*, in *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, STOC ’05*, pp. 481–485 (Association for Computing Machinery, New York, NY, USA) (2005).
- [361] R. Van Meter, *Quantum networking*, chapter Distributed Digital Computation and Communication, pp. 126–128 (John Wiley & Sons, Ltd) (2014).
- [362] L. Wagner and S. McDonald, *Finding traitors in secure networks using Byzantine agreements*, [International Journal of Network Security](#) **10**, 11 (2010).
- [363] S. Gaertner, M. Bourennane, C. Kurtsiefer, A. Cabello, and H. Weinfurter, *Experimental demonstration of a quantum protocol for Byzantine agreement and liar detection*, [Physical Review Letters](#) **100**, 070504 (2008).
- [364] M. Marozzi and L. Mostarda, *Quantum consensus: An overview*, [arXiv:2101.04192](#) (2021).
- [365] 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).

- [366] H. J. Metcalf and P. van der Straten, *Laser cooling and trapping* (Springer New York) (2001).
- [367] 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).
- [368] J. Borregaard, M. Zugenmaier, J. M. Petersen, H. Shen, G. Vasilakis, K. Jensen, E. S. Polzik, and A. S. Sørensen, *Scalable photonic network architecture based on motional averaging in room temperature gas*, [Nature Communications](#) **7** (2016).
- [369] C. B. Møller, R. A. Thomas, G. Vasilakis, E. Zeuthen, Y. Tsaturyan, M. Balabas, K. Jensen, A. Schliesser, K. Hammerer, and E. S. Polzik, *Quantum back-action-evading measurement of motion in a negative mass reference frame*, [Nature](#) **547**, 191 (2017).
- [370] K. Shinbrough, B. D. Hunt, and V. O. Lorenz, *Optimization of broadband  $\Lambda$ -type quantum memory using Gaussian pulses*, [Physical Review A](#) **103**, 062418 (2021).
- [371] E. L. Wooten, K. M. Kissa, A. Yi-Yan, E. J. Murphy, D. A. Lafaw, P. F. Hallemeier, D. Maack, D. V. Attanasio, D. J. Fritz, G. J. McBrien, and D. E. Bossi, *A review of lithium niobate modulators for fiber-optic communications systems*, [IEEE Journal of Selected Topics in Quantum Electronics](#) **6**, 69 (2000).
- [372] D. Zhu, L. Shao, M. Yu, R. Cheng, B. Desiatov, C. J. Xin, Y. Hu, J. Holzgrafe, S. Ghosh, A. Shams-Ansari, E. Puma, N. Sinclair, C. Reimer, M. Zhang, and M. Lončar, *Integrated photonics on thin-film lithium niobate*, [Advances in Optics and Photonics](#) **13**, 242 (2021).
- [373] 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).
- [374] C. E. Rogers and P. L. Gould, *Nanosecond pulse shaping at 780 nm with fiber-based electro-optical modulators and a double-pass tapered amplifier*, [Optics Express](#) **24**, 2596 (2016).
- [375] E. Gehrig and O. Hess, *Optimization of tapered semiconductor optical amplifiers for picosecond pulse amplification*, [Applied Physics Letters](#) **87**, 251106 (2005).
- [376] A. F. Forrest, M. Krakowski, P. Bardella, and M. A. Cataluna, *Double-pass amplification of picosecond pulses with a tapered semiconductor amplifier*, [Optics Express](#) **27**, 30752 (2019).
- [377] 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).



- [378] I. H. Malitson, *Interspecimen comparison of the refractive index of fused silica*, [Journal of the Optical Society of America](#) **55**, 1205 (1965).
- [379] T. Peters, B. Wittrock, F. Blatt, T. Halfmann, and L. P. Yatsenko, *Thermometry of ultracold atoms by electromagnetically induced transparency*, [Physical Review A](#) **85**, 063416 (2012).
- [380] M. W. Sørensen and A. S. Sørensen, *Three-dimensional theory for light-matter interaction*, [Physical Review A](#) **77**, 013826 (2008).
- [381] 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).
- [382] 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).
- [383] R. Mottola, *Room temperature single-photon sources and atomic quantum memories for broadband quantum networks*, phdthesis, Universität Basel, Klingelbergstrasse 82, 4056 Basel (2023).
- [384] 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).
- [385] 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 narrow-band photon pairs at the  $^{87}\text{Rb}$  D1 line*, [Optics Express](#) **28**, 3159 (2020).
- [386] 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).
- [387] Z. Y. Ou and Y. J. Lu, *Cavity enhanced spontaneous parametric down-conversion for the prolongation of correlation time between conjugate photons*, [Physical Review Letters](#) **83**, 2556 (1999).
- [388] 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).
- [389] 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).

- [390] 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).
- [391] 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).
- [392] D. M. Aikens, *The truth about scratch and dig*, in *International Optical Design Conference and Optical Fabrication and Testing*, p. OTuA2 (Optical Society of America) (2010).
- [393] 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).
- [394] Y. Wang, W. Yang, Z. Li, and Y. Zheng, *Determination of blue-light-induced infrared absorption based on mode-matching efficiency in an optical parametric oscillator*, [Scientific Reports](#) **7** (2017).
- [395] W. Nagourney, *Quantum electronics for atomic physics*, chapter Energy relations in optical cavities, p. 41 (Oxford University Press), first edition (2010).
- [396] 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).
- [397] 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).
- [398] B. E. A. Saleh and M. C. Teich, *Fundamentals of photonics*, (John Wiley & Sons, Inc.) (1991).
- [399] 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).
- [400] S. Bettelli, *Comment on “Coherence measures for heralded single-photon sources”*, [Physical Review A](#) **81**, 037801 (2010).
- [401] 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).
- [402] 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).
- [403] K. Zhang, J. Guo, L. Q. Chen, C. Yuan, Z. Y. Ou, and W. Zhang, *Suppression of the four-wave-mixing background noise in a quantum memory retrieval process by channel blocking*, [Physical Review A](#) **90**, 033823 (2014).

- [404] I. Vurgaftman and M. Bashkansky, *Suppressing four-wave mixing in warm-atomic-vapor quantum memory*, [Physical Review A](#) **87**, 063836 (2013).
- [405] N. Allard and J. Kielkopf, *The effect of neutral nonresonant collisions on atomic spectral lines*, [Reviews of Modern Physics](#) **54**, 1103 (1982).
- [406] 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).
- [407] 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).
- [408] B. S. Fong, M. Davies, and P. Deschamps, *Timing resolution and time walk in super low  $K$  factor single-photon avalanche diode—measurement and optimization*, [Journal of Nanophotonics](#) **12**, 1 (2018).
- [409] Technical brief 76W-28079-0, *Using Tektronix AWG7000 series in synchronous applications*, [techreport, Tektronix](#) (2011).
- [410] Highland Technology, *T564 digital delay generator*, 650 Potrero Avenue, San Francisco, CA 94110 (2019).
- [411] 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).
- [412] I. E. 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).
- [413] G. Ghosh, *Dispersion-equation coefficients for the refractive index and birefringence of calcite and quartz crystals*, [Optics Communications](#) **163**, 95 (1999).
- [414] A. W. Ziarkash, S. K. Joshi, M. Stipčević, and R. Ursin, *Comparative study of afterpulsing behavior and models in single photon counting avalanche photo diode detectors*, [Scientific Reports](#) **8** (2018).
- [415] 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).
- [416] J. Keaveney, C. S. Adams, and I. G. Hughes, *ElecSus: Extension to arbitrary geometry magneto-optics*, [Computer Physics Communications](#) **224**, 311 (2018).

- [417] L. Weller, *Absolute absorption and dispersion in a thermal Rb vapour at high densities and high magnetic fields*, [phdthesis](#), Durham University (2013).
- [418] A. B. Matsko, I. Novikova, M. O. Scully, and G. R. Welch, *Radiation trapping in coherent media*, [Physical Review Letters](#) **87**, 133601 (2001).
- [419] M. Klein, Y. Xiao, M. Hohensee, D. F. Phillips, and R. L. Walsworth, *Slow light and EIT under realistic (imperfect) conditions*, in S. M. Shahriar, P. R. Hemmer, and J. R. Lowell (editors), *Advances in Slow and Fast Light II*, volume 7226, pp. 102–109, [International Society for Optics and Photonics \(SPIE\)](#) (2009).
- [420] J. Kitching, *Chip-scale atomic devices*, [Applied Physics Reviews](#) **5**, 031302 (2018).
- [421] E. Ramsay, *Solid immersion lens applications for nanophotonic devices*, [Journal of Nanophotonics](#) **2**, 021854 (2008).
- [422] 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).