


ORCA

Towards an integrated noise-free Quantum Memory

KRZYSZTOF TADEUSZ KACZMAREK
LADY MARGARET HALL, OXFORD



A thesis submitted for the degree of
Doctor of Philosophy
Trinity term 2017

ORCA - towards an integrated noise-free Quantum Memory

KRZYSZTOF TADEUSZ KACZMAREK

LADY MARGARET HALL, OXFORD

Submitted for the degree of Doctor of Philosophy, Trinity term 2017

Abstract

Light is an ideal information carrier for quantum networks: its quantum properties are not degraded by noise in ambient conditions and it has a large information capacity owing to its high bandwidth. However, quantum technologies based on photonic networks have been hampered by the exponentially poor scaling of the underlying probabilistic quantum operations. Quantum optical memories, devices that store, manipulate, and release on-demand quantum light, have been identified as an indispensable network component, because they facilitate scalability. Noise-free operation of the memory is critical, since even small additional noise can render the memory classical by destroying the quantum character of the light.

Here I introduce a new broadband quantum memory protocol - the off-resonant cascaded absorption (ORCA) memory - that is inherently noise-free and operates in ambient conditions. I model ORCA theoretically, after which I characterise the classical performance of a proof-of-concept implementation in warm caesium vapour, using weak near-infrared coherent states. In order to verify quantum operation of the ORCA memory, I interface it with a broadband heralded single-photon source. I measure and compare the quantum statistics of the stored and retrieved light, observing for the first time their full preservation in a room-temperature atomic quantum memory. Finally, I make headway into integrating the ORCA vapour memory by investigating an alkali-filled hollow-core fibre platform. I use light-induced atomic desorption (LIAD) to demonstrate record-breaking alkali vapour densities in fibre, a prerequisite for efficient memory operation.

Because of its technical simplicity and integrability, its high bandwidth and its low noise, ORCA provides a viable route towards next generation, photonic quantum network technologies.

This thesis is dedicated to Maxwell's

Acknowledgments

I could not have hoped for a more enriching PhD experience than what I received in Oxford and the Ultrafast Group. As usual in life, these 4 years had its ups and downs, but they all provided an invaluable lesson in what it means to be an independent researcher. For example, how a dead-end project can become an impulse to invent a new type of quantum memory!

I feel very fortunate to have met the people I encountered in my time in Oxford, all of whom could not possibly be listed in such a small space. I would therefore like to start by thanking the big boss himself, #yaboyIAW. Ian, it has been a privilege to be part of such an exciting moment in Ultrafast's history. Josh, dude, I've said this many times, but thank you again so much for your help over the years. I honestly couldn't ask for a better supervisor. Good luck in Bath! Dylan, we've been in this memory business together from the very start and you've truly grown it into something amazing. Keep kicking-ass! Michael, sorry for ending the project that you started, but we did get a paper out of it! Steve, thanks for providing the critical impulse to start the experiment. Best of luck in London! Patrick (L.), mate, thanks for blatantly doing a lot of the work described in this thesis. Academia needs more people as talented and good-natured as you! Maka guts the bantz will stop. Amir, thanks for all your help and advice. Please pass on my regards to our machine overlords. Eilon, thanks for being such an enthusiastic supporter of crazy science. Helen, thanks so much for the office banter. Hope some of that "global warming" will make those Berlin winters bearable. Benni, my only hope is that with a kid on your back you will finally stop being such a perfect scientist! Patrick M., thanks for setting the bar for hard work at an impossible high, forcing me to find my own work

standard.

Joe, thanks for being a constant reminder of what it means to be a good student. Though I think in Oscar you will find a worthy competitor in computational prowess and keen attention to detail. Sarah, another perfect student! Don't worry about the rest of your PhD, you have a couple theses' worth of material already. Shaobo, you have quite big footsteps to follow, but with your remarkable rate of growth, I am sure you will overtake them in no time. Guillaume, you got some truly nice results in the few months we worked together. Keep up the good work! Tom, nicely done on those filters man, hopefully you'll get a chance to use them! James, good work with really completing your experiment from start to finish, you should be proud! Joelle, well done on taking over OxOPS, and good luck with writing up! Finally, big shout-out to everyone in the Ultrafast Group and everyone who left. Will definitely miss Pubfest and punting with you guys and gals!

Now that we are on the topic of drinking however, I could not avoid extending the hugest thanks to the folks at the LMH MCR. It's your fault this PhD took so long! But it's also your fault I now have friends for life in almost every corner of the planet. At this point it is unfortunately impossible to mention everyone by name, but you know who you are. And thank you for that. Also, special thanks to Gio Maschio. Without Qigong I would have probably lost it already in year 2!

Finally, I owe a debt of gratitude to my parents. Without them, I would literally not be here. But also, without their consequent inspiration and advice, I would probably not end up with a PhD in Quantum Physics from the University of Oxford.

Now to do it justice...

Author's publications

Publications based on material in thesis

1. **Kaczmarek, K. T.**, Ledingham, P. M., Brecht, B., Thomas, S. E., Thekkadath, G. S., Lazo-Arjona, O., Munns, J. H. D., Poem, E., Feizpour, A., Saunders, D. J., Nunn, J., Walmsley, I. A. (2017). *A high-speed noise-free optical quantum memory*. arXiv preprint: 1704.00013 [quant-ph].
2. **Kaczmarek, K. T.**, Saunders, D. J., Sprague, M. R., Kolthammer, W. S., Feizpour, A., Ledingham, P. M., Brecht, B., Poem, E., Walmsley, I. A., Nunn, J. (2015). *Ultra-high and persistent optical depths of caesium in Kagomé-type hollow-core photonic crystal fibres*. Optics Letters, 40(23).

Related publications

3. Nunn, J., Munns, J. H. D., Thomas, S., **Kaczmarek, K. T.**, Qiu, C., Feizpour, A., Poem E., Brecht, B., Saunders, D. J., Ledingham, P. M., Reddy, D. V., Raymer, M. G. Walmsley, I. A. (2017). *Theory of noise suppression in Λ -type quantum memories by means of a cavity*. Physical Review A, 96(1).
4. Thomas, S. E., Munns, J. H. D., **Kaczmarek, K. T.**, Qiu, C., Brecht, B., Feizpour, A., Ledingham, P. M., Walmsley, I. A., Nunn, J., Saunders, D. J. (2017). *High Efficiency Raman Memory by Suppressing Radiation Trapping*. New Journal of Physics, 19(6).
5. Saunders, D. J., Munns, J. H. D., Champion, T. F. M., Qiu, C., **Kaczmarek, K. T.**, Poem, E., Ledingham, P. M., Walmsley, I. A., Nunn, J. (2015). *A Cavity-Enhanced Room-Temperature Broadband Raman Memory*. Physical Review Letters, 116(9).
6. Michelberger, P. S., Champion, T. F. M., Sprague, M. R., **Kaczmarek, K. T.**, Barbieri, M., Jin, X. M., England, D. G., Kolthammer, W. S., Saunders, D. J., Nunn, J., Walmsley, I. A. (2015). *Interfacing GHz-bandwidth heralded single photons with a room-temperature Raman quantum memory*. New Journal of Physics, 17(4).

Other publications

7. Kiffner, M., Feizpour, A., **Kaczmarek, K. T.**, Jaksch, D., Nunn, J. (2016). *Two-way interconversion of millimeter-wave and optical fields in Rydberg gases*. New Journal of Physics, 18(9).
8. Poem, E., Weinzetl, C., Klatzow, J., **Kaczmarek, K. T.**, Munns, J. H. D., Champion, T. F. M., Saunders, D. J., Nunn, J., Walmsley, I. A. (2015). *Broadband noise-free optical quantum memory with neutral nitrogen-vacancy centers in diamond*. Physical Review B, 91(20).

Contents

Contents	x
List of Figures	xiv
1 Introduction	1
1.1 Figures of merit	3
1.1.1 Efficiency	3
1.1.2 Lifetime	4
1.1.3 Storage bandwidth	4
1.1.4 Time-bandwidth product	5
1.1.5 Noise level	5
1.1.6 Memory fidelity	6
1.1.7 Modal capacity	6
1.1.8 Technical simplicity	7
1.2 Quantum memory applications	7
1.2.1 Quantum storage	8
1.2.2 Quantum communication	9
1.2.3 Buffering	11
1.2.4 Multiplexing	12
1.2.5 Quantum state manipulation	15
1.2.6 Tests of fundamental physics	17
1.3 Review of quantum memories	18

1.3.1	Echo-based memories	19
1.3.2	Non-echo memories	22
1.4	Introducing ORCA	25
1.5	Outline of thesis	27
2	Theory of the Off-Resonant Cascaded Absorption (ORCA) memory protocol	28
2.1	Equations of motion	29
2.1.1	Atomic part	30
2.1.2	Optical part	31
2.1.3	Interaction part	32
2.1.4	Optical Bloch equations	33
2.1.5	Maxwell's equations	36
2.1.6	Maxwell-Bloch equations	37
2.2	Analytic solution	38
2.2.1	Short cell limit	39
2.2.2	Long cell limit	41
2.3	Numerical solution	45
2.3.1	Comparison with analytics	45
2.3.2	Numerical analysis	47
3	Classical characterisation of the ORCA memory in warm caesium vapour	55
3.1	Experimental setup	56
3.1.1	Memory cell	57
3.1.2	Signal	59
3.1.3	Control	61
3.1.4	Detection	62
3.1.5	Evidence of ORCA operation	63
3.2	Memory lifetime	64
3.2.1	ORCA lifetime theory	65
3.2.2	ORCA lifetime measurement	66

3.3	Memory efficiency	67
3.4	Memory noise level	69
4	Interfacing the ORCA memory with a single-photon source	72
4.1	Single-photon source	73
4.1.1	Parametric downconversion theory	73
4.1.2	Photon statistics	75
4.1.3	Nonlinear crystal	78
4.1.4	Crystal mount	80
4.1.5	Experimental setup	81
4.2	Interfacing PDC source with ORCA in Cs	83
4.2.1	Data collection	84
4.2.2	Results	88
5	Cs-filled hollow-core photonic crystal fibres as an integrated platform for quantum memories	95
5.1	Cs-filled hollow-core fibres	96
5.1.1	Kagomé-structured hollow-core fibres	97
5.1.2	Filling the HC-PCFs with Cs	99
5.2	Light-Induced Atomic Desorption	101
5.2.1	Generating ultra-high optical depths of Cs in HC-PCFs using NIR LIAD	102
5.2.2	Modelling the NIR LIAD experiment	107
5.2.3	Generating even higher optical depths using green LIAD	112
6	Conclusion	118
6.1	Summary of thesis	119
6.2	Outlook	121
6.2.1	Investigating the modal capacity of ORCA	122
6.2.2	Improving the ORCA memory performance	123
6.2.3	Verifying ORCA quantum operation	125
6.2.4	Improving ORCA integration	126

<i>CONTENTS</i>	xiii
References	130
Appendix A. Optical depth	148
Appendix B. Laser locking	152
B.1 Locking the ECDL reference	153
B.2 Locking the signal Ti:Sapph	154
B.3 Locking the control Ti:Sapph	154
Appendix C. $g^{(2)}$	155
Appendix D. Polarisation	156

List of Figures

1.1	Quantum repeater scheme	10
1.2	N-photon coincidence waiting times	14
1.3	Raman memory and FWM noise	24
2.1	ORCA memory scheme	29
2.2	Comparison of ORCA analytics and numerics	46
2.3	Typical ORCA numerical trace (weak coupling)	49
2.4	ORCA efficiency as a function of Δ and n (simulation)	50
2.5	ORCA efficiency as a function of $\Delta\tau$ and $E_{\text{in/out}}$ (simulation)	51
2.6	Typical ORCA numerical trace (strong coupling)	52
3.1	ORCA experimental setup	56
3.2	ORCA level scheme in Cs	57
3.3	Signal Ti:Sapph spectrum	60
3.4	Control Ti:Sapph spectrum	61
3.5	Schematic of measurement settings	63
3.6	Typical ORCA memory time trace (experiment)	64
3.7	ORCA memory lifetime curve	66
3.8	ORCA efficiency vs control pulse energy (experiment)	68
3.9	ORCA noise measurement	71
4.1	Scheme of measuring the Glauber correlation functions	76
4.2	PDC spectra	78

4.3	PDC phasematching curves	79
4.4	PDC mount schematic	80
4.5	Experimental scheme for the heralded single-photon source	81
4.6	PDC filter spectrum	82
4.7	PDC singles traces	86
4.8	PDC doubles and triples traces	87
4.9	PDC ORCA efficiency over time	89
4.10	ORCA $g^{(1,1)}$ over time	91
4.11	$g_h^{(2)}$ measurement	93
5.1	Kagomé fibre characterisation	98
5.2	Optical images of the kagomé fibres	98
5.3	Ultra-high vacuum system scheme	100
5.4	Loading fibres with Cs	101
5.5	NIR LIAD experimental setup	103
5.6	NIR LIAD desorbing laser scheme	104
5.7	LIAD in 26 μm fibre	105
5.8	LIAD in 46 μm fibre	106
5.9	Pulsed NIR LIAD (experiment)	107
5.10	LIAD model	108
5.11	Long-term NIR LIAD model	110
5.12	Pulsed NIR LIAD model	111
5.13	Green LIAD experimental setup	114
5.14	Low power green LIAD	115
5.15	High power green LIAD	116
A.1	Cs absorption spectra	151
B.1	CW reference locking scheme	152
B.2	Locking signal for CW reference	153
D.1	ORCA efficiency versus control polarisation	157

D.2 Polarisation-sensitive ORCA lifetime measurement 158

1

Introduction

Quantum networks, structures which process and transmit information encoded in the quantised degrees of freedom of the information carrier, hold great potential for implementing quantum technologies, which promise to revolutionise computing, communication and metrology [1]. This has motivated a tremendous research effort into different architectures for implementing these networks [2, 3, 4]. Light is an ideal information carrier for quantum networks: its quantum properties are not degraded by noise in ambient conditions and it has a large information capacity owing to its high bandwidth [5]. Furthermore, light is especially well-suited for integration [6], decreasing the footprint of a photonics-based technology. However, quantum technologies based on photonic networks to-date have been hampered by the exponentially poor scaling of the underlying probabilistic quantum operations [7]. One identified solution to this scalability problem has been quantum optical memories [8].

A quantum memory¹ in principle is a device that allows one to store and retrieve on-demand² a quantum optical “signal”, i.e. a quantum state of light, such as a single photon. In more general terms, a quantum optical memory essentially performs a unitary beamsplitter-type transformation between the moving optical signal mode and a stationary matter mode³ [12, 13]. This is quite different from a classical memory, which would operate by measuring a quantum state and then re-producing it (so-called measure-and-prepare strategy), as in that case, storage would lead to collapse of the state’s wavefunction and loss of coherence. This can be seen as a result of the quantum no-cloning theorem [14].

In this thesis, I demonstrate a new type of quantum memory based on two-photon off-resonant cascaded absorption (ORCA) in warm vapour. Similarly to many other memory protocols (see section 1.3), it operates by manipulating the absorption of an ensemble of atoms using a strong auxiliary, or “control”, optical field. The light-matter interaction strength can be increased by integrating the memory within a waveguide structure, such as a photonic crystal fibre [15], potentially also decreasing the device’s footprint. The interaction of a strong field with a large ensemble of atoms however, produces noise in many memory protocols (especially at room-temperature), contaminating and deteriorating the stored quantum state of light [16]. Therefore, it is critical for any quantum memory application (see section 1.2), to verify the noise-free performance of the underlying memory.

I quantify the noise generated by a memory and describe other popular metrics used to evaluate these devices in section 1.1 of this chapter. This is followed by an overview of the important applications of quantum memories (section 1.2). In section 1.3 I review the current state-of-the-art of quantum memory research and applications. Finally, I introduce the new ORCA memory protocol I conceived in section 1.4, and conclude the chapter with an outline of the thesis.

¹I will be using the terms quantum memory and quantum optical memory interchangeably, though the former is more general and in principle includes devices, which can be prepared in a long-lived quantum state, without actually absorbing a quantum state of light [9, 10].

²By on-demand, we mean it is possible to dynamically, i.e. during storage, choose when to retrieve the stored photon.

³Unless the memory is based on switchable optical cavities [11], as mentioned in section 1.3, in which case the “matter” mode is a localised (in a cavity) optical mode.

1.1 Figures of merit

Before discussing the various applications and examples of quantum memories, it is useful to first define the various parameters with which to judge them (Table 1.1). I will define these figures of merit as pertaining to on-demand, ensemble-based memories, though many of them are equally applicable to single-emitter systems (see section 1.3). Of course, aside from the quantitative parameters described below, it is beneficial, especially for the scalability of memory-enhanced photonic networks, that the quantum memory is technically simple (ideally operable in ambient conditions), integratable and robust. For specific applications, it may also be desirable that the memory is capable of storing light of a particular wavelength, e.g. in one of the telecommunications bands (e.g. $1.5\ \mu\text{m}$) for quantum communication (section 1.2.2).

Efficiency	η
Lifetime	$t_{1/e}$
Bandwidth	δ
Time-bandwidth product	B
Noise level	μ_1
Fidelity	$\mathcal{F}$

Table 1.1: A summary of the figures of merit for quantum memories, as described in this thesis.

1.1.1 Efficiency

I define the total memory efficiency η as the amount of signal light retrieved from the memory compared to the amount that was sent in⁴. We can furthermore distinguish between the storage or read-in efficiency η_{in} , i.e. how much light is absorbed into the memory medium, and the retrieval or read-out efficiency η_{out} , i.e. how much of that light is then retrieved from it. The total memory efficiency is then the product of the two, i.e. $\eta = \eta_{\text{in}}\eta_{\text{out}}$. In many memory protocols, the non-stored light will be transmitted through the memory. This allows us to treat the memory interaction as a light-matter beam splitter, where η_{in} and η_{out} can be seen as the reflectivities of the beam splitters representing the storage and retrieval interactions [12].

⁴For example, one can compare the energy contained in the stored and retrieved light pulses, or the probabilities of having a single photon stored and retrieved.

For practical applications, we are usually more interested in the end-to-end efficiency of the memory η_{e2e} . This not only includes the internal memory efficiency η , but also the losses introduced by other components necessary to operate the memory, such as filters, vapour cell walls, fibre couplings, i.e. $\eta_{e2e} = \eta_T \eta$ (η_T quantifying the transmission of the memory system). Usually, η_T can be improved with proper engineering and using better quality components, while η may have more fundamental limitations⁵, yet for many practical applications both efficiencies have to be high (see section 1.2).

1.1.2 Lifetime

In on-demand quantum memories we can dynamically control the amount of time we want to store a light pulse for, or storage time. Once light is stored as an atomic coherence however, the coherence itself will deteriorate over time due to various dephasing and loss processes (see section 3.2), lowering the retrieval efficiency. This will limit the maximum storage time of the memory. We can therefore define the memory lifetime $t_{1/e}$ as the storage time after which the retrieval efficiency drops to $1/e$ of its original value.

1.1.3 Storage bandwidth

A memory's storage bandwidth is defined as the maximum bandwidth, or spectral width, of a light pulse that can be stored in the memory. The storage bandwidth therefore determines how temporally short pulses of light can be stored in the memory. For transform-limited pulses, the bandwidth of a pulse δ is related to its duration τ by the simple expression:

$$\delta\tau = F, \tag{1.1}$$

where F is the Fourier constant, whose magnitude depends on the shape of the pulse, e.g. $F \approx 0.44$ for a Gaussian-shaped pulse (as we often assume for the output of a mode-locked laser). Assuming one can produce a train of pulses separated by their duration, the storage bandwidth will determine the maximum clock rate over which a memory can operate.

⁵For example, re-absorption of the light as it propagates through the storage medium [17].

1.1.4 Time-bandwidth product

As the name suggests, the time-bandwidth product B^6 is the product of the memory lifetime $t_{1/e}$ and bandwidth δ , and is closely related to the maximum number of clock cycles over which a memory can operate before it dephases, i.e. $B = t_{1/e}\delta \approx \frac{t_{1/e}}{\tau}$. It is worth mentioning though, that if one wishes to operate a memory in feed-forward or conditional operation, the practical time-bandwidth product will often be limited by the bandwidth of the feed-forward electronics ($\sim$ GHz [18]) rather than the optical bandwidth, and so one should rather consider the product of the experimental clock rate R and storage time $B' = Rt_{1/e}$.

1.1.5 Noise level

All of the figures of merit described above are equally important for classical optical memories as much as for quantum memories. However, classical memories, often operating on coherent states with very large photon number, will be less sensitive to any noise processes present in the memory. In the case of storing a single photon in a quantum memory, any “noise” photons produced by the memory interaction itself could contaminate the signal, precluding the use of the memory for quantum applications. The amount of noise produced by different quantum memories can be nicely compared by the so-called μ_1 value [19], defined as:

$$\mu_1 = \langle N^{\text{noise}} \rangle / \eta \sim \frac{1}{\text{SNR}}, \quad (1.2)$$

i.e. the average number of noise photons produced during a storage/retrieval attempt $\langle N^{\text{noise}} \rangle$, normalised by the memory efficiency η . For an input signal of average photon number $\langle N_{\text{in}} \rangle = 1$, the μ_1 will be equal to the inverse of the signal-to-noise ratio (SNR) at the memory output, which is another common metric for noise [15].

However, it is insufficient to predict guaranteed quantum operation by only measuring the mean of the noise (which the μ_1 quantifies), because even a very small average amount of noise can impair quantum signatures if the variance of the noise is large, e.g. thermal [16]. Quantum-limited operation can be verified by measuring the normalised Glauber

⁶Also known as delay-bandwidth product or fractional delay.

correlation functions (described in more detail in section 4.1.2) [20], as changes in this quantity will signify altered photon-number statistics⁷.

1.1.6 Memory fidelity

Somewhat related to the noise level and memory efficiency described in the previous sections is the memory fidelity, which can be defined as the amount of overlap between the input state of light described by density matrix ρ_{in} and the retrieved state ρ_{out} , i.e.

$$\mathcal{F} = \text{tr} \left(\sqrt{\sqrt{\rho_{\text{in}}} \rho_{\text{out}} \sqrt{\rho_{\text{in}}}} \right)^2. \quad (1.3)$$

An inefficient or noisy memory will have a reduced fidelity in the photon number basis, but the way we define fidelity also takes into account any changes induced by the memory on the mode of the retrieved light, e.g. by spectral-temporal or polarisation distortion⁸.

For a memory to be truly quantum, its fidelity (in the degree of freedom we wish to encode information in) has to be better than the maximum fidelity achievable with a classical measure-and-prepare strategy, i.e. 50% for a single photon (limited by quantum no-cloning) [21]. For practical quantum information processing applications however, the fidelity has to be high enough such that fault-tolerant quantum error correction methods can overcome memory imperfection [22].

1.1.7 Modal capacity

As has been mentioned before, one can think of an on-demand quantum memory as a beamsplitter coupling an optical mode and a matter mode. Conventional optical beam-splitters are multimode in time and frequency, i.e. they will efficiently operate on many different time-frequency modes. Quantum memories however, especially those based on atomic transitions with narrow homogenous linewidths, will often only store a certain number of optical modes. One can then define the modal capacity of a memory in terms

⁷It is important to note that in certain cases, e.g. Rydberg-based memories, one may wish to modify the quantum statistics of the stored light, as explained in section 1.2.5.

⁸Of course, as we will see in section 1.2.5, it is sometimes desirable to change the quantum state upon retrieval, in which case fidelity in that specific degree of freedom ceases to be a good metric, and it may be beneficial to consider instead the overlap of the output state with a desired state.

of how many modes, and in what degree of freedom, can a memory store efficiently⁹. One way to theoretically quantify this for a memory protocol is to perform a singular-value decomposition of the Green's function describing the linear memory interaction, which will provide the eigenmodes and associated efficiencies (obtainable from the eigenvalues) of the memory [23]. The memory's modal capacity can be related to the amount of quantum information it can store simultaneously [24].

1.1.8 Technical simplicity

As mentioned in section 1.1, another important aspect of a quantum memory, though harder to quantify¹⁰, is the ease of implementation or integration. For many of the applications outlined in the next section, a number of quantum memories is required. This means a potential memory platform has to be scalable, i.e. it has to be easy to produce in large numbers. This often means that a memory has to be operable at the quantum level ideally at close to ambient conditions (room temperature and pressure)¹¹.

Ultimately, as described at the start of this chapter, our goal is to implement a quantum memory in a photonic quantum network. These are often based on optical fibres [25] or integrated optical circuits [6]. Therefore it is desirable that a practical quantum memory is compatible with these platforms, which means the memory itself is either fibre-based [15] or integrated[26].

1.2 Quantum memory applications

The basic application of a quantum memory is to store quantum information, like a hard-drive for a quantum computer. By storing information, one can also buffer it, in order to synchronise the different components of a quantum computer or network, which take different amounts of time to operate. However, if we look at an atomic quantum memory as a light-matter interface, rather than a simple storage device, the range of applications grows significantly. This basically arises from the fact that by storing an optical field as

⁹Efficient in this context has been defined before as e.g. 70% [23].

¹⁰One possible way to quantify this is by the cost of producing a single quantum memory unit.

¹¹Though with the current progress in cryogenics and atom-chip systems it is possible that low-temperature memories will become scalable.

a material excitation, it allows one to manipulate and measure the light by manipulating and measuring the properties of the atoms. This enables certain operations on light, which would not be easily performed directly.

In the following sections I outline the different applications of quantum memories, and point towards the memory parameters (as defined in section 1.1) critical for each particular application.

1.2.1 Quantum storage

As a fundamentally non-stationary object, photons are excellent at transmitting information. However, in order to store that information, it has to be translated into a stationary medium¹². The basic purpose of a quantum memory is therefore to faithfully preserve quantum information encoded in light. For this application, the crucial figure of merit will be memory fidelity (ideally as close to 100%, or at least above the threshold given by quantum error correction), which will depend on memory efficiency, lifetime and noise level. Multimodality could also lead to a higher information capacity, e.g. a spatially multimode memory could be used to store images [29].

Lifetime-wise, the precise maximum desired storage time of the memory will depend on the application one has in mind. For quantum information processing applications (see section 1.2.3), if there is a need to store the result of some quantum operation while waiting for another part of the system to finish its work, then the timescales, at least for photonics, would be $\sim \mu\text{s}$ [11]. However, if one wishes to store traditionally long-term classical information (such as genomic data, etc.) secured by a quantum key, quantum key distribution schemes can be employed [30], which benefit from relatively long-lived ($\sim \text{h}$ [31]) memories, as explained in the next section.

¹²One can imagine using optical fibre loops in order to hold light-encoded quantum information [27], however one does not have control then on when to release it. Using switches solves this problem, yet introduces more loss [28].

1.2.2 Quantum communication

Quantum communication, in particular quantum key distribution¹³, is generally based on similar principles as the quantum no-cloning theorem, i.e. any attempt at intercepting the key and extracting information from it will measurably degrade it, which can be picked up by the receiver [33]. In 1991 Ekert discovered that one can use shared entanglement between the parties in order to transmit a quantum-secure random key [34].

While light is by far the best way to transmit quantum information over long distances, owing to its speed and insensitivity to noise in ambient conditions, it also faces a limit on maximum transmission distance, stemming from losses in optical fibre¹⁴, which are on the order of 0.2 dB/km [36]. This means that after 100 km, the signal amplitude would fall to 1% of its initial value. If the carrier of our signal is a single photon, which cannot be deterministically amplified without adding noise (as it would violate the quantum no-cloning theorem [14]), channel loss will reduce the key sharing rates [25].

Instead of sending the carrier of quantum information directly, one can imagine dividing the transmission link into smaller, less lossy segments, over which one shares entanglement. One then also requires a way to swap entanglement between the segments, in order to send the information over the whole transmission link, as shown in Fig. 1.1 A and B. This is the basis of a quantum repeater¹⁵, originally proposed in 1998 [37]. The quantum repeater thus overcomes the exponential scaling of signal transmission error with distance [25], with only a polynomial overhead in time and a logarithmic overhead in the number of nodes [37].

¹³Once two parties share a secure key, then there are provably secure classical means of encrypting information (such as a message) using that key. Hence for secure communication, the parties only have to share the key securely, while the encrypted message can be sent over a non-secure channel [32].

¹⁴Recently, quantum information has been successfully transmitted using free space beam propagation between satellites and on-earth receiver stations [35].

¹⁵Similarly, a classical telecom link will be divided into smaller segments, interspersed with classical repeaters based on amplifiers, since classical communication is less sensitive to noise.




Figure 1.1: **A, B**, quantum repeater based on entanglement swapping. An entangled state is produced and a Bell-type measurements between the photons from neighbouring link segments are used to spread entanglement over the whole transmission link. **C, D**, entanglement swapping using DLCZ-based quantum repeaters. An entangled state between an atomic ensemble and photon is produced, and a Bell-type measurement of the photons from neighbouring ensembles leads to entangling of the two ensembles. Retrieving photons from the ensembles, one can then perform a similar operation leading to entanglement spreading over the whole transmission link.

In principle, a combination of sources of entangled states and Bell-type measurements is all that is required to transmit entanglement. However, in most cases, the generation, transmission and measurement of the state will be probabilistic, and thus the probability of successful swapping of entanglement over the whole transmission link will scale poorly with the number of nodes, i.e. with transmission distance. This is where quantum memories enter, as a means to synchronise the successful entanglement swap operations. One can imagine generating two spatially-separated entangled pairs of photons, using one

of each pair to swap entanglement, and once successful, store the other of each pair in a quantum memory. One can then wait for other link segments to be established in similar fashion, and then release the photons in order to join the links together, reducing the requirement for all link segments to be established simultaneously.

One of the most well-known quantum repeater protocols is the DLCZ protocol (from Duan, Lukin, Cirac and Zoller, authors of the original proposal) [38], in which the entangled photon pair source is integrated with the quantum memory, as shown in Fig. 1.1 C and D. Spontaneous Raman scattering [39] is used to generate an entangled state of light and atomic ensemble. Interfering and detecting the light from two DLCZ memories can then be used to generate entanglement between the ensembles [40]. Another Raman scattering process can then be used to map the state of the ensemble into an optical field, which can then be used to establish further entangled segments.

A useful quantum repeater requires sufficient storage time and efficiency¹⁶, such that the probability of establishing an entangled link is higher than the success probability of sending a photon directly through a lossy fibre [41]. Multimode capacity and large memory bandwidth can also be a strong asset for this application, since they can improve the rate at which entanglement swapping operations can be performed [42].

1.2.3 Buffering

Traditionally, it was believed photons do not interact strongly with one another¹⁷, and thus implementing quantum information processing with them using linear operations (which is the basis of linear optical quantum computing, or LOQC [47]) would be difficult, due to the difficulty of constructing conditional (e.g. two-qubit) gates.

In 2001 Knill, Laflamme and Milburn demonstrated that efficient quantum computation is possible using only single-photon sources, linear optical networks and photon-detectors¹⁸ [49]. The KLM scheme improves the scalability of probabilistic two-qubit

¹⁶Storage and retrieval efficiency in the case of an “in/out” quantum memory (as are the focus of this thesis) and retrieval efficiency in the case of the DLCZ memory.

¹⁷There has recently been enormous progress in generating strong single-photon nonlinearities [43], by (linear) coupling of light to highly nonlinear media, such as single emitters (single trapped atoms [44], colour defects in diamond [45]) or atomic ensembles in Rydberg states [46].

¹⁸Since then, there have been significant progress in the field of LOQC, with a significant amount of work being currently devoted to one-way or cluster-state computation, in which one first creates a very large and complex entangled state, and then performs computations by specific measurements on the

gates by harnessing quantum teleportation [50]. The key to this scheme is the realisation that one can perform the non-deterministic gate operation independently many times, and once successful, teleport the two qubits one wishes to manipulate onto the output of the gate, thus performing the gate operation on them. While the gate operation is being attempted on an auxiliary system, the quantum state of the qubit is stored in a quantum memory. Buffering is naturally related to the multiplexing application described in the next section, and requires memories with similar characteristics.

More generally, quantum memories can be used as buffers of quantum information, enabling the synchronisation of probabilistic operations in different parts of a quantum photonic circuit, and facilitating feed-forward operation of conditional processing steps. For these applications, the time-bandwidth product is paramount, as it determines the number of computational clock cycles over which a memory can buffer. Of course, in order to reduce errors in the computation, it is also crucial that the memory preserves the qubit well, placing stringent requirements on memory fidelity, noise level and efficiency. Since two-photon gates in LOQC depend on the interference, and thus indistinguishability, of the photons, single-mode memories may be preferred for this application, as they provide a natural mode filter for the photons.

1.2.4 Multiplexing

I touched upon temporal multiplexing in the previous two sections by mentioning synchronisation of probabilistic operations. In general, multiplexing (as defined here) involves switching the output of many concurrent probabilistic processes to a single mode¹⁹. If that is possible, then one can construct a near-deterministic process by selecting on the successful outcome of many trials of probabilistic processes. Here I will focus on multiplexing in the time domain using quantum memories, though there is considerable research effort being conducted in multiplexing spectrally using nonlinear processes [52] or spatially, using electro-optic switches [53].

consituent systems [48].

¹⁹This is slightly different from the usual definition of multiplexing, as e.g. used in telecommunications, where multiplexing broadly means switching many input modes into a single mode, in order to increase the information capacity of the latter. For example in telecom wavelength-division multiplexing, many frequency modes are reduced to a single spatial mode, after which they are de-multiplexed and can be operated on independently [51].

One well-investigated example is multiplexing of probabilistic heralded photon sources [54, 55], which historically have been the workhorse of quantum optics. Heralded photon sources based on parametric processes, such as downconversion or four-wave mixing, generate pairs of photons (more on this in section 4.1.1). Detecting one photon then heralds the presence of another. This pair generation process occurs probabilistically, with the success probability q kept necessarily low²⁰, in order to reduce the probability of generating more than one pair (again, more on this in section 4.1.1).

Quantum memories can be used to increase q by temporal multiplexing, i.e. by storing the photon when a pair was generated successfully (as evidenced by detecting the herald) in a chosen time window, and then releasing it at the end of the window. This approach reduces the maximum photon generation rate by the number of clock cycles or time bins over which we wish to multiplex, since we end up using only one of the time bins in the chosen time window²¹, though it does improve the photon generation rate in the final time bin. Multiplexing over $1/q$ time bins, the probability of generating a photon will then be limited by the memory efficiency, or in other words, the photon generation probability is given by $\sim \eta Bq \approx \eta$. The source/memory combination becomes akin then to a “deterministic” quantum dot source, with a non-unit collection efficiency [57].

Similarly to the example of quantum communication, another approach is to integrate the source and memory, as in DLCZ. Here, aside from temporal multiplexing, another interesting approach is to use spatial multiplexing, where one heralds not only on the existence of a DLCZ-generated photon, but also on its position. This is because DLCZ is spatially multimode, since the initial photons can be generated in any direction, and its momentum/position will be entangled with the atomic ensemble, and thus the momentum/position of the retrieved photons [58]. One can then modify the angle of the retrieved photons depending on the detected herald position, in order to retrieve them all in a single mode. As with temporal multiplexing, the probability of generating a photon will be low in each spatial mode, but the probability of generating it in any spatial mode may approach unity [59].

²⁰The highest q which can be achieved using a photon source based on two-mode squeezing is 25%, assuming no loss and photon-number resolving detection of the herald [56].

²¹Though, assuming 100% memories, the actual total photon generation rate does not change.

It would be difficult for a single memory-enhanced non-deterministic photon source to outperform an actual “deterministic” source (based e.g. on quantum dots [57]) in single photon production rates. However, the advantage of using even inefficient quantum memories becomes apparent, when we consider the generation of N -photon states. The probability of N independent probabilistic sources simultaneously producing photons will scale as q^N . The blue diamonds in Fig. 1.2 show how the average time required to generate a single N -photon state scales with N , assuming we attempt to generate 1 GHz bandwidth pairs in each source pumped at a rate of 1 GHz with a probability $q = 0.013$ ²², compared with a similar waiting time from a spatially de-multiplexed quantum dot source, with a photon extraction efficiency of $q_{\text{dot}} = 0.28$ [60]. It is clearly visible that the N -photon production rate scales much more favourably with N for a single de-multiplexed deterministic source than for N probabilistic sources.




Figure 1.2: Comparison of average waiting times required to produce a single N -photon state, between state-of-the-art de-multiplexed quantum dots, and memory-synchronised and non-synchronised probabilistic photon sources based on PDC (see text for more detail).

²²The pair generation probability used in a paper by Kaneda et al. demonstrating temporal multiplexing of two PDC sources using a quantum memory based on free-space cavities [11].

Now let us imagine inserting a quantum memory after each probabilistic heralded photon source. Whenever any of the N sources produces a photon, we store in the memory. Once all N sources have produced a photon, we release the N photons from our memories all at once. It turns out, this greatly improves the N -photon generation rate [55], as shown in Fig. 1.2, which assumes 50% efficient memories with a 100 ns lifetime. This improvement can be roughly understood as stemming from the earlier observation that now the total system success probability would scale as something closer to $\sim (\eta Bq)^N$ (at least in the limit $\eta Bq < 1$) [55, 11].

It is immediately apparent that for temporal multiplexing applications an important figure of merit is the time-bandwidth product, adjusted by the memory efficiency. However, less explicit, but equally important is the noise level, as deteriorating the single-photon characteristics of the memory output will render the memory useless for multiplexing single-photon sources altogether. While in the case of a spatially multiplexed DLCZ memory, the protocol benefits from high spatial multimodality, interfacing parametric photon sources with a single-mode memory may improve the purity or indistinguishability of generated photons, by “cleaning up” their mode.

1.2.5 Quantum state manipulation

As the name suggests, LOQC is based on linear optics, which are in general also passive or time stationary objects. This limits the degree of manipulation of photonic states that is possible with them²³. A light-matter interface or quantum memory can increase our ability to manipulate quantum states of light, by introducing time non-stationary or nonlinear operations, either in the process of coupling of light to atoms or in the atoms themselves.

First of all, as mentioned before, many quantum memory protocols can be reduced to a unitary and linear light-matter beamsplitter interaction²⁴. Similar to a classical beamsplitter, which couples and thus allows manipulation of different spatial modes, a quantum memory thus allows coupling of different time bin modes²⁵ [63, 12]. This, combined with

²³Which is also why LOQC requires such a significant resource overhead [7].

²⁴The analogy is strong enough that one can observe memory-mediated Hong-Ou-Mandel interference [61].

²⁵Not to be confused with temporal modes [62], which will be discussed later.

the DLCZ-enabled two-mode squeezing operation, was shown by Humphreys et al. to be sufficient for generating and manipulating continuous-variable cluster states, and consequently quantum computation [64].

The key to the Humphreys scheme lies in the fact that in many quantum memory protocols (described in section 1.3), the coupling of light to matter is mediated by an auxiliary optical field, which in principle can be time non-stationary, especially when it is pulsed. This makes it possible for the memory to be mode-selective, i.e. enable coupling only of a specific spectral-temporal mode of light to the material system [62]. By appropriately choosing the auxiliary field for read-out, one can then couple the material mode to a different optical mode. A quantum memory thus allows one to manipulate the spectral-temporal properties of the optical field [65, 66], which is impossible to achieve with time stationary elements.

The applications outlined above still rely on the quantum memory performing linear operations on light. As has been mentioned before, light, especially on the single-photon level, was historically considered to be incapable of experiencing strong nonlinearities due to its weak, non-resonant interaction with the environment and itself. The highly efficient storage of light using quantum memories already disproves the notion that light cannot be made to interact strongly with matter. Furthermore, quantum memories can be based on material systems that are highly nonlinear in themselves, such as single atoms [44], or highly-excited or “Rydberg” ensembles of atoms [67].

Storing light in Rydberg ensembles and manipulating the atoms using microwave fields resonant with the transitions between Rydberg states²⁶ would enable manipulation of the statistics of the stored light [70]. Furthermore, due to the Rydberg blockade effect²⁷, a Rydberg-based memory can enable strong nonlinearities akin to a single atom [72]. The nonlinearity comes from strong coupling of light to an easily (single-photon level) saturable medium.

Rydberg ensembles have been used to demonstrate cross-phase modulation of single-

²⁶Rydberg atoms coupled to optical fields and microwaves have also been used for microwave-to-optical conversion [68, 69].

²⁷Strong dipole-dipole or van der Waals interactions, that arise for highly excited atoms, lead to shifts in the resonances of neighbouring atoms, and thus prevent on-resonance excitation of multiple atoms within a certain radius, i.e. the Rydberg blockade radius [71]. Within the blockade radius, the atomic ensemble thus behaves like a single Rydberg-excited “super-atom” [67].

photon level pulses [73] and even single-photon level pulse collisions [74]. The former is useful for constructing controlled phase gates, as required for quantum computing [75]. Furthermore, Rydberg ensembles have also been used for single-photon level self-phase modulation [76], with potential applications in single-photon filters or sources.

For quantum state manipulation, it is important that the memory does not destroy the quantum properties of the stored light, i.e. that it does not introduce noise. As has been mentioned, in the case of photon state manipulation, fidelity, as defined in section 1.1.6, ceases to be an appropriate metric. Lifetime is also less important for this application, however it may be beneficial for the memory to have a large acceptance bandwidth. Single-modeness is also crucial, since we wish to selectively couple and manipulate specific spectral-temporal modes.

1.2.6 Tests of fundamental physics

As we have seen in the previous sections, quantum memories are a critical component of quantum communication and information processing architectures. Aside from the practical applications of such quantum devices, they also allow us to investigate and test the limits of quantum theory, e.g. memory-enhanced quantum communication enabling tests of non-locality over large distances [77], or light-matter beamsplitters exhibiting Hong-Ou-Mandel interference [61].

In general quantum memories have been shown to be an interesting tool to investigate more fundamental aspects of quantum theory. Examples being generating [78, 79], quantifying [80] and verifying macroscopic entanglement [81, 82]. Interestingly, storing a single photon in an atomic ensemble entangles the macroscopic number of constituent atoms [83]. It is then in principle possible to independently measure the state of the atoms [84], allowing one to potentially directly measure this large-scale entanglement. Another interesting experiment involved generating and measuring Einstein-Podolsky-Rosen position-momentum correlated states using a spatially multimode DLCZ-type memory [58].

In order to investigate quantum theory in great detail, it is most important that the memory does not introduce noise, which would make the retrieved state classical. The

other figures of merit are less important here than they were for the practical applications outlined in the previous sections, though memory efficiency and lifetime may still be critical in certain cases, e.g. for closing loopholes in tests of non-locality [77]. In order to store entanglement, the memory benefits from being multimode in the entangled degrees of freedom [85]. Otherwise, the modes can be stored separately in independent single-mode memories [86].

1.3 Review of quantum memories

In the last twenty years there has been considerable research effort in developing quantum memories for the applications outlined in the previous section. Here, I review the current state-of-the-art, and briefly describe the main memory protocols and storage media currently being investigated. While there are many ways in which one could divide the different existing memory protocols²⁸, I have chosen to broadly classify them in two categories, based on whether they rely on controlled rephasing techniques for storage and retrieval. While one can imagine storing light in an optical fibre [27] or atom-based [89] delay line, without the ability to extract the light on-demand, many of the applications outlined in the previous sections become impossible.

Worth noting is the recent progress in using electro-optically switchable optical cavities, in order to implement an on-demand quantum memory [11, 90]. The benefits of this approach are inherently noise-free operation (since there are no other optical fields or atoms involved that could generate noise) and the ability to store ultrashort ($\sim$ fs) pulses at various wavelengths (since these memories do not rely on narrow atomic resonances). However, the operating speed of the memory will depend on the bandwidth of the electro-optic switch²⁹ ($\sim$ GHz at best [18]) and the cavity/loop round-trip time. Decreasing the latter though will decrease the memory lifetime (limited by propagation loss in these systems), since the most lossy component will in most cases be the switch [7]. Still, temporal multiplexing of two PDC sources has been successfully demonstrated using a switchable free-space cavity memory [11].

²⁸The presented selection is by no means exhaustive, and does not include e.g. continuous-variable memories [87] or certain more exotic echo-based protocols [88].

²⁹Note the distinction between B and B' in section 1.1.4.

There has also been recent substantial progress in investigating single emitters for quantum memory applications, e.g. single atoms trapped inside optical cavities [91] or colour centres in diamond, such as nitrogen-vacancy [92] or silicon-vacancy centres [93]. As has been mentioned in section 1.2.5, single emitter systems inherently offer very strong nonlinear interactions and can be used also as sources of single photons [45]. The problem with their practical application is the inherently low coupling strength between a single photon and single emitter, which necessitates using either optical cavities [91] or nanophotonic structures [45] to increase the interaction strength, either by making the photon interact with the emitter many times (the former) or focusing the photon down to a very small region around the emitter (the latter). Nanophotonic single-emitter systems furthermore give the promise of scalability and integrability with other integrated photonic architectures [94].

One can increase the light-matter coupling strength by using atomic ensembles, since by significantly increasing the number of atoms involved in the interaction, one increases the probability that a photon will interact with one of them [95]. However, since on-demand quantum memories will predominantly require a strong auxiliary field to mediate the transformation of a photon into a matter excitation, the interaction of this field with the large number of atoms (either via spontaneous scattering or more involved phenomena) will often lead to the generation of noise. Some of which is not easily filtered from the desired single-photon level signal, preventing true quantum operation [16].

Regardless, ensemble-based memories have received significant attention over the years, and many of them have already reached the performance levels required for many of the applications outlined in section 1.2. Moreover, ensemble-based memories tend to be technically simpler than single-emitter systems, and many of them can operate in ambient conditions. In the following I will limit the discussion to ensemble-based on-demand quantum memories.

1.3.1 Echo-based memories

An inhomogeneously broadened ensemble of two-level atoms will absorb and emit light. The light, which we henceforth call the “signal”, will be stored as a superposition of coher-

ences generated in each frequency class of atoms. The phases of these coherences will then evolve at different rates, depending on the frequency class, leading to dephasing of the total coherence³⁰. In “naturally” broadened media, this dephasing will be irreversible. However, one can imagine inducing controlled and reversible broadening of an atomic ensemble using external electric or magnetic fields. This is the basis of the Controlled Reversible Inhomogenous Broadening (CRIB) [96] and Gradient Echo Memory (GEM) [97] protocols. In the former, a gradient field is introduced transversally to the light propagation direction, while in the latter it is applied longitudinally. This gradient field shifts the frequencies of the ensemble, generating strong additional inhomogenous broadening. After some time, the gradient can be reversed, inverting the phase accumulation rate in the different frequency classes of atoms, allowing them to rephase. Once they do, this leads to re-emission of the stored light via a superradiance-like effect [96].

The storage time of GEM and CRIB can be significantly extended by sending in a strong control field, which transfers the stored coherence from the short-lived excited state to a third meta-stable state³¹ [98]. The atomic states in this case are usually arranged in a so-called Λ -type level configuration, where the “storage” state lies energetically below the excited state. The initial and storage states may be e.g. two different spin states of a ground state manifold [99]. In that case, in order for the memory to operate, it is first necessary to prepare the atomic ensemble in one of the two states and empty the other, e.g. via optical pumping [100]. This will be common for most Λ -type memory schemes.

CRIB has originally been implemented in a solid-state system, a rare-earth-ion doped crystal (REID) [101]. These systems offer very long lifetimes, even in the excited state, due to a shielding effect of their optically-addressable valence electrons by the other electrons in outer shells [31]. However, the ions still suffer from coupling to the vibrations of the host crystal lattice, i.e. phonons, and thus must be cryogenically cooled [102]. While the ion density can be quite high in REIDs, the electric dipole moments of the optical transitions are usually quite low (which is related to the long lifetimes of the excited states

³⁰We will come across this again in section 3.2.

³¹Actually, the control field dresses the excited state and then adiabatically transfers the stored coherence to the meta-stable state, similarly to the Raman memory (albeit at smaller detunings) described in the next section. In the dressed state picture, the three-level protocol behaves as the original (quasi-)two-level one described earlier.

[103]), weakening the light-ion coupling [104].

Another limitation to the CRIB memory efficiency arises because of re-absorption of the signal during retrieval, as it propagates through the memory medium. This is overcome in GEM, in which (due to the longitudinally applied broadening field gradient) the signal will interact with a certain frequency class of atoms in only a very thin slice of the memory medium [97]. GEM in laser-cooled atomic ensembles has reached efficiencies of 87% and millisecond lifetimes [41], and storage fidelities of 98% for weak coherent states in warm ensembles [105]. GEM also has very good multimode capabilities [106], with photon time bin manipulation been demonstrated [63]. Storage of single photons has not yet been achieved (even though the noise level was measured to be low [41]), and the bandwidth of GEM has so far been limited to $\sim\text{MHz}$ ³² [108].

Both GEM and CRIB require an external field gradient, which needs to be inverted at some point, in order to rephase the stored coherences generated in the inhomogeneously broadened ensemble. This can be understood as arising from the Fourier transform, i.e. if the absorption spectrum of the ensemble has no periodicity, then one expects no periodicity in the emission in the time domain, or in other words, no spontaneous rephasing. However, one can imagine generating, e.g. via optical pumping, a periodic structure in an inhomogeneously broadened ensemble, in order to produce periodic spontaneous emission. This is the basis of the Atomic Frequency Comb (AFC) memory protocol [109, 89]. An absorptive comb structure is generated optically in an inhomogeneously broadened ensemble. After absorbing light, the stored coherences will dephase and rephase with a period given by the inverse of the comb spacing. Similarly to the other echo-based protocols, rephasing of the coherence will lead to re-emission of the stored light.

While the traditional two-level AFC has to be prepared prior to storage, which also pre-defines the storage time, a Λ -type variant of the protocol, called spin-wave AFC, uses a control field to freeze the rephasing process and transfer the excitation to a spin coherence [110]. One can then dynamically choose when to apply the control field again

³²In principle, for a pure three-level system the bandwidth of GEM should only be limited by the achievable field gradient [107]. However, a higher gradient requires higher density of atoms, in order to keep the number of atoms per frequency bin the same, for good light-matter coupling. Also, most atomic systems will have a more complicated electronic structure over the scale of hundreds of megahertz, and thus storing pulses of comparable bandwidth would be difficult due to interference between the different transitions.

in order to transfer the coherence back to the excited state and re-launch the rephasing process, granting on-demand memory operation.

While the two-level AFC in REIDs is quite established, with nanophotonic integration [26], cavity-enhanced high-efficiency storage [111], broadband telecom photons [110] and hyper-entanglement [112] been demonstrated, and the protocol having been used for information processing [65] and investigations of macroscopic entanglement [82], there has been recent progress in improving the performance of the spin-wave AFC, with 12% efficiency [113] and storage of single photons been demonstrated [114]. There have even been demonstrations of an AFC analogue of DLCZ in REIDs [115, 42]. The acceptance bandwidths of spin-wave AFC have been quite low ($\sim$ MHz) however and the protocol currently requires time-consuming preparation of the atomic ensemble prior to storage, significantly reducing operating speeds [116].

1.3.2 Non-echo memories

The on-demand memory protocols described in the previous section rely on the absorption of light by an ensemble of two-level atoms with appropriately engineered inhomogenous broadening, combined with or followed by transfer of the stored coherence to a third meta-stable state. It is however possible to directly map an optical field to a spin coherence, without engineering or populating the excited state, via a Raman-like transition (remember DLCZ in section 1.2.2).

In fact, the first proposed memory protocols were based on such a process [117, 118]. The Electromagnetically-Induced Transparency (EIT) or slow-light memory relies on using an auxiliary optical “control” field to modify the dispersion of an atomic medium, in order to compress and slow the signal (to be stored) down to a halt within the medium [119, 120]. This can be understood as the control field dressing the excited state of the ensemble, creating a transparency window within the atomic absorption spectrum³³ [121]. This reduction in absorption will be, via Kramers-Kröning relations, accompanied by a strong dispersive feature, slowing and compressing the signal field. By adiabatically turning the control field off, the transparency window disappears, absorbing the resonant

³³This transparency arises from destructive interference between the two generated dressed states.

signal field into a spin coherence. Turning the control field on again reverses the process, re-accelerating the signal.

EIT has been implemented in warm [122] and cold [120] alkali ensembles, as well as in REIDs [123]. Storage times of minutes in cold atoms [124], and just under a minute in REIDs [125], have been shown, along with memory efficiencies of 78% [126] (76% [127]) in the former (latter). Optical state manipulation has also been demonstrated [128]. Single photon storage in cold [129] and warm [130] alkali vapours has been attempted using DLCZ-based photon sources. The storage bandwidth of EIT is determined by the size of the transparency window, which depends on the Rabi frequency of the control field and the optical depth [131] (related to the number density of atoms, see [Appendix A. Optical depth](#)). However, EIT has so far been limited to $\sim$ MHz bandwidths, since it relies on precise quantum interference in a three-level system, whereas the excited states of alkalis are composed of a manifold of states separated by tens of MHz, which may contribute additional excitation paths deteriorating the quality of interference.

Much larger, $\sim$ GHz and $\sim$ THz storage bandwidths have been demonstrated in warm alkali and hydrogen vapours [17, 132], and diamond [133], respectively, using the off-resonant Raman memory protocol [134], leading to memory time-bandwidth products on the order of thousands. Here one actually uses the absorption feature generated by a strong off-resonant control field to absorb the signal and transfer it to a spin coherence, via a STIRAP³⁴-like process. The bandwidth of the Raman memory is thus determined by the bandwidth of the control pulse³⁵. In fact, the Raman memory, for low coupling strengths, is single-mode and thus mode-selective [136], which means an appropriately shaped control field can be used to filter out an appropriately matched signal from a superposition of spectral-temporal modes [62] or manipulate the signal upon read-out [66]. Because of its off-resonant character (and thus very little linear absorption, when the memory is off), the system can be reduced to a beamsplitter-like interaction between different travelling optical modes and a stationary matter mode [12, 137], which has been used to demonstrate Hong-Ou-Mandel interference [61]. Polarisation qubit storage in a

³⁴Stimulated Raman Adiabatic Passage [135], where two strong fields are used to move population from one atomic state to another via a stimulated two-photon Raman transition.

³⁵In the absence of selection rules, the bandwidth will be limited eventually by the ground-storage state separation (few GHz in alkalis).

dual-rail configuration ³⁶ has been demonstrated in a Raman memory in warm vapour [86], as was polarisation entanglement storage in a cold ensemble [138]. Integration of the memory using alkali-filled hollow-core fibres has also been demonstrated [15].

Single-photon storage has been demonstrated in the Raman memory, both in room-temperature diamond [133], and warm [16] and cold [138] alkali vapours, with non-classical photon number statistics being observed upon retrieval in diamond and cold atoms. However, in all cases the photon number statistics were modified in the memory due to the presence of noise. Four-wave mixing (FWM), due to the control field coupling to the signal transition and generating DLCZ-type coherence between the spin states (depicted in Fig. fig:fig1d2), has been identified as a major source of noise in off-resonant Λ -type memories [16, 139]. Since FWM, like DLCZ, generates a two-mode squeezed state of so-called Stokes (in the signal mode) and anti-Stokes photons, if one filters out the anti-Stokes, the Stokes photons will have thermal statistics [42], strongly disrupting the photon number statistics in the signal mode. Furthermore, the Raman memory relies on very strong filtering of the signal/Stokes mode, in order to separate it from the much stronger ($\sim 10^9$ times stronger for single-photon storage) control³⁷ [16].




Figure 1.3: On the green background is shown a schematic of Stokes Raman scattering, which is the desired process behind the Raman memory (signal is green, control is brown). On the blue background is shown the simultaneous undesired process of anti-Stokes scattering (anti-stokes field is blue), which leads to four-wave mixing noise (brown background) in the signal field (green arrow). Figure adapted from [140].

³⁶Because of polarisation selection rules, the Raman memory in alkalis is single-mode in polarisation as well.

³⁷This also limits the Raman memory's end-to-end efficiency.

At high coupling strengths, FWM will also lead to gain, limiting the achievable memory efficiency for high-fidelity storage [100]. FWM was shown to be reduced in a cavity³⁸-enhanced Raman memory [140] or through introducing loss via absorption in the anti-Stokes mode [141]. A cavity furthermore should impose single-mode memory operation even at high memory efficiencies (by keeping the coupling strength per cavity roundtrip low) [139].

1.4 Introducing ORCA

To-date most on-demand memory protocols have been based on spin storage in Λ -type level schemes. This has been motivated by the desire to extend the memory lifetime, historically with a quantum repeater application in mind. However, as we have seen (especially in room-temperature systems), the Λ level scheme is vulnerable to a number of sources of noise, e.g. control-field induced fluorescence [142] or four-wave mixing [16], preventing actual quantum operation.

If one relaxes the requirement on memory lifetime (which is not critical to many of the applications outlined in section 1.2), one is then free to investigate non-spin states for storage. One is in fact, not even limited to Λ -type level schemes. For example, highly excited Rydberg states are capable of spontaneous emission lifetimes of tens of microseconds [67]. In fact, narrowband ultracold EIT schemes, based on storage in a Rydberg state via a two-photon ladder or cascade transition, have been used to store classical light and generate non-classical states on output [70], as well as store single photons generated from another cold atomic ensemble [143]. Room-temperature storage using a similar scheme would be difficult however³⁹, since the difference in wavelength between the two fields (usually many nanometres) leads to strong Doppler broadening of the two-photon transition⁴⁰, and thus dephasing of the memory⁴¹. Using low-lying (e.g. orbital) states, which have a shorter (yet still sufficient) lifetime, and where the two

³⁸Resonant with the Stokes and control modes, but anti-resonant with the anti-Stokes.

³⁹Not impossible though, since one can use a three-photon transition with a non-collinear beam geometry, in order to eliminate Doppler broadening [144].

⁴⁰In Λ -type schemes, where the signal/control wavelength difference is on the order of GHz, this is much less of an effect, though still measurable [145].

⁴¹See section 3.2 for a more detailed explanation.

transitions are much closer in wavelength, thus gives the promise of room-temperature operation.

This observation led me to conceive a new quantum memory protocol, based on off-resonant cascaded absorption (ORCA) in a three-level ladder, or cascade, configuration that is optimised for warm atomic vapours [146]. The operational principle is similar to the Raman memory (see section 1.3.2): two broadband optical fields, signal and control, satisfy a two-photon resonance condition with a storage state, while individually being far-off resonance from their respective single-photon transitions (see chapter 2 for more details). Unlike in Λ -type protocols, the ORCA memory bandwidth is not in principle limited by the ground state splitting [138], and a counter-propagating beam geometry reduces Doppler broadening, and thus motional-induced dephasing of the stored coherence, giving the potential for functional time-bandwidth products [18].

In the new scheme, the storage state is a doubly-excited electronic state, which has no thermal excitations even at room temperature. Therefore the protocol in principle requires no preparation of the atomic ensemble prior to storage (increasing the operating speed of the memory), and there is no contamination of the retrieved fields due to imperfect optical pumping. This points to the major benefit of ORCA in that it is *noise-free*. The signal and control wavelengths can be chosen such that the control photons are significantly detuned from the populated ground-state transition (THz detunings are readily available in the rich atomic structure of alkalis [147]). This effectively eliminates any control-field induced fluorescence noise [148]. More importantly though, due to the ladder configuration, there is no scattering process that could populate the storage state, and so four-wave mixing noise is eliminated. Finally, efficient suppression of control field leakage on the output detection is readily achievable using off-the-shelf low-loss interference filters, and bolstered by the counter-propagating field configuration, in principle enabling *end-to-end* device efficiencies approaching the *internal* memory efficiency [18]. The ORCA scheme is therefore a broadband and noise-free quantum optical memory protocol operating within a technically simple, room-temperature platform.

1.5 Outline of thesis

In chapter 2, I describe the theory behind ORCA, first deriving the equations of motion that govern the memory interaction, and then presenting analytical solutions to these equations, in certain parameter regimes. This is followed by results of numerically solving these equations. Using these results, I estimate the achievable ORCA memory performance.

In the next chapter (3), I present the experimental implementation of ORCA in warm caesium vapour, and its characterisation using broadband (~ 1 GHz bandwidth) weak coherent states. I find the memory to behave as predicted by the theory, with no measurable noise.

In chapter 4, I interface the ORCA memory from the previous chapter with a broadband single-photon source based on parametric downconversion in a nonlinear crystal waveguide. I verify noise-free operation of the memory by directly showing that the photon number statistics of the signal remain unchanged by storage and retrieval. Part of the material presented in chapters 3 and 4 has been published as an online preprint in [146].

Chapter 5 is devoted to an investigation of a potential integrated platform for the ORCA memory, i.e. a caesium-filled hollow-core photonic crystal fibre. I find that the ambient atomic densities in the fibre are too low for memory operation, but I can generate record-high transient densities using the phenomenon of light-induced atomic desorption (LIAD), which I model in our system. Part of the material presented in this chapter has been published in [95].

Finally, in chapter 6, I present an outlook of the potential future experiments which could be made on ORCA, and discuss ways of improving the memory performance.

2

Theory of the Off-Resonant Cascaded Absorption (ORCA) memory protocol

In this chapter I will derive the simplified Maxwell-Bloch equations that describe the ORCA (Off-Resonant Cascaded Absorption) memory [146] interaction. For the most part the derivation is similar to the thoroughly investigated Raman memory in three-level Λ -type atomic systems (following [23] and [149]). The main difference, aside from the energy level configuration, is that in the Raman memory both optical fields are co-propagating, while in ORCA the fields are made to be counter-propagating to reduce motion-induced (or Doppler) dephasing (see section 3.2.1). This makes the ORCA equations less straightforward to solve analytically.

I will nonetheless try to find an analytical solution for certain limiting cases, following the approach of master student Guillaume Thekkadath, who worked on this alongside

me. I will then briefly compare these analytical results with numerical solutions to the ORCA equations of motion, using a code for solving simplified Maxwell-Bloch equations, originally written for a different project [150]. Finally, I will investigate and optimise the performance of ORCA numerically in more detail using a more comprehensive programme written in Python by DPhil student Oscar Lazo-Arjona (based on code originally written in Matlab by myself and Dr. Amir Feizpour for solving the full Maxwell-Bloch equations in three-level systems).

2.1 Equations of motion




Figure 2.1: The ORCA memory protocol. **A**, Storage: a weak “signal” field (blue) and strong “control” (orange) are overlapped counter-propagating in an atomic vapour. The broadband fields are on two-photon resonance with a doubly-excited storage state $|3\rangle$, while remaining far-detuned from the intermediate state $|2\rangle$. **B**, The storage process maps the signal field to an atomic coherence or “orbital wave” (dashed line) between the ground $|1\rangle$ and storage $|3\rangle$ states. **C**, Recall: re-applying the control pulse after the desired storage time leads to re-mapping of the atomic coherence back into an optical field and thus re-emission of the signal in the forward direction.

The Off-Resonant Cascaded Absorption memory protocol involves three atomic energy levels: ground $|1\rangle$, intermediate $|2\rangle$ and storage $|3\rangle$ in a ladder or cascade configuration; and two off-resonant counter-propagating optical fields: the weak “signal” we wish to absorb on the $|1\rangle \rightarrow |2\rangle$ transition and a strong “control” on the $|2\rangle \rightarrow |3\rangle$ (as shown in Fig. 2.1). We use the control to map the propagating optical signal into a stationary atomic coherence between states $|1\rangle$ and $|3\rangle$ in a large atomic ensemble, which I call an *orbital wave*¹, and then map it back into an optical signal after a dynamically chosen storage time. The storage process is based on off-resonant two-photon absorption, while retrieval can be seen as stimulated two-photon emission.

The Hamiltonian describing the light-matter system outlined above can be written as

$$\hat{H} = \hat{H}_A + \hat{H}_L + \hat{H}_{ED}, \quad (2.1)$$

where $\hat{H}_A$ is the Hamiltonian for the atom, $\hat{H}_L$ is the optical free-field Hamiltonian and $\hat{H}_{ED}$ is the electric dipole Hamiltonian, responsible for the light-matter interaction. In the next sections, I will describe each of these terms in more detail, after which I will derive the equations of motions for the atomic states in the Heisenberg picture. Then I will use Maxwell’s equations to describe the propagation of the optical fields, and finally combine the two sets of equations, known jointly as the Maxwell-Bloch equations.

2.1.1 Atomic part

The Hamiltonian for the bare atomic system can be described in the basis of its three eigenstates $|1\rangle$, $|2\rangle$, $|3\rangle$:

$$\hat{H}_A = \sum_j \hbar\omega_j |j\rangle\langle j| = \sum_j \hbar\omega_j \hat{\sigma}_{jj}, \quad (2.2)$$

where ω_j is the resonant frequency of the state $|j\rangle$. Since $|j\rangle$ are eigenstates, they are orthonormal, i.e. $\langle i|j\rangle = \delta_{ij}$ (δ_{ij} being the Kronecker delta), and thus form a natural basis to describe the atomic states and their dynamics. In fact, the electric dipole operator $\mathbf{d}$ can be written as

¹Analogously to the term *spin wave* used in Λ -type memories [110, 17], where the coherence is created between electronic spin states.

$$\mathbf{d} = \sum_{j,k} \mathbf{d}_{jk} \hat{\sigma}_{jk}. \quad (2.3)$$

When $j \neq k$, we call $\hat{\sigma}_{jk} = |j\rangle\langle k|$ “flip” (or transition projection) operators as they will change the state of the atom, i.e. $\hat{\sigma}_{jk}|k\rangle = |j\rangle$. The coefficients $\mathbf{d}_{jk} = \langle j|\mathbf{d}|k\rangle$ are the dipole matrix elements.

2.1.2 Optical part

The total electric field in a three-level memory is composed of two components, a strong classical control field $\mathbf{E}_c$, and a weak quantum signal field $\mathbf{E}_s$, which we wish to store:

$$\mathbf{E} = \mathbf{E}_c + \mathbf{E}_s \quad (2.4)$$

For the rest of this chapter, I will only consider the fields in one dimension, both propagating along the z -axis and ignoring their transverse behaviour, i.e. in the paraxial approximation [23]. The classical control field can then be written as:

$$\mathbf{E}_c = \mathbf{v}_c E_c(z, t) e^{i\omega_c(t+z/c)} + \text{c.c.}, \quad (2.5)$$

where $\mathbf{v}_c$ is the control polarisation vector, ω_c is the central frequency of the control field, and $E_c(t, z)$ is the slowly varying envelope of the control, which describes its spatio-temporal shape.

The signal field, in principle a single photon, requires a quantum description and thus I write it as:

$$\mathbf{E}_s = i\mathbf{v}_s \int g(\omega) \hat{a}(\omega) e^{-i\omega z/c} d\omega + \text{h.c.}, \quad (2.6)$$

where $\mathbf{v}_s$ is the signal polarisation vector, $g(\omega) = \sqrt{\hbar\omega/4\pi\epsilon_0 c}$ is the mode amplitude of the signal and $\hat{a}(\omega)$ is the narrowband annihilation operator for a photon with frequency ω [151]. We now make the transformation $\hat{a}(\omega, t) = \hat{a}(\omega) e^{i\omega t}$. Assuming the bandwidth of the signal pulse is small compared to its central frequency ω_s , I assume $g(\omega) \approx g(\omega_s) = g_s$ and re-write the signal as

$$\mathbf{E}_s = i\mathbf{v}_s g_s S(z, t) e^{i\omega_s(t-z/c)} + \text{h.c.}, \quad (2.7)$$

where S is the slowly varying broadband annihilation operator:

$$S(z, t) = \frac{e^{-i\omega_s(t-z/c)}}{\sqrt{2\pi}} \int d\omega \hat{a}(\omega, t) e^{-i\omega z/c}. \quad (2.8)$$

The way I defined $\mathbf{E}_c$ and $\mathbf{E}_s$, the signal and control fields are moving in opposite directions along the z -axis. The free-field Hamiltonian for light can be written as a sum of the Hamiltonians for the classical field $\hat{H}_C$ and quantum field $\hat{H}_Q$:

$$\hat{H}_L = \hat{H}_C + \hat{H}_Q = \frac{\epsilon_0}{2} \int dt (\mathbf{E}_c^2 + c^2 \mathbf{B}_c^2) + \int d\omega \hbar \omega (\hat{a}(\omega)^\dagger \hat{a}(\omega) + 1/2), \quad (2.9)$$

where $\mathbf{B}_c$ is the magnetic field associated with the electromagnetic wave defined by the classical electric field $\mathbf{E}_c$.

2.1.3 Interaction part

The atom-field interaction in the dipole approximation² is given by:

$$H_{ED} = -\mathbf{d} \cdot \mathbf{E}, \quad (2.10)$$

where the dipole operator $\mathbf{d}$, defined in terms of the atomic electron position $\mathbf{r}$, is

$$\mathbf{d} = -e\mathbf{r}, \quad (2.11)$$

e being the elementary charge [152]. The electric dipole interaction has odd parity, which means only the off-diagonal elements of the dipole operator are non-zero, i.e. $\mathbf{d}_{jk} = 0$ for $j = k$ [23].

Finally, I assume there is no direct dipole transition between states $|1\rangle$ and $|3\rangle$, leaving me with:

²The dipole approximation is justified in this case, since the wavelengths of the light inducing atomic transitions (~ 900 nm) are much larger than the size of the atom (~ 0.1 nm).

$$H_{ED} = -(\mathbf{d}_{12}\hat{\sigma}_{12} + \mathbf{d}_{23}\hat{\sigma}_{23} + \text{h.c.})\mathbf{E}. \quad (2.12)$$

2.1.4 Optical Bloch equations

In order to describe the memory interaction, we are interested in how the atomic states evolve in time. To answer this question I use the Heisenberg equation of motion, i.e.

$$\partial_t \hat{\sigma}_{jk} = i[\hat{\sigma}_{jk}, \hat{H}] = i[\hat{\sigma}_{jk}, \hat{H}_A + \hat{H}_{ED}], \quad (2.13)$$

since the free-field hamiltonian for light $\hat{H}_L$ commutes with the atomic operators $\hat{\sigma}_{jk}$. I therefore get the following set of equations:

$$\begin{aligned} \partial_t \hat{\sigma}_{11} &= -i(\mathbf{d}_{12}\hat{\sigma}_{12} - \text{h.c.})\mathbf{E}, \\ \partial_t \hat{\sigma}_{33} &= -i(\mathbf{d}_{23}\hat{\sigma}_{23} - \text{h.c.})\mathbf{E} \end{aligned} \quad (2.14)$$

for the atomic populations³, and:

$$\begin{aligned} \partial_t \hat{\sigma}_{12} &= i\omega_{21}\hat{\sigma}_{12} - i[\mathbf{d}_{12}^*(\hat{\sigma}_{11} - \hat{\sigma}_{22}) + \mathbf{d}_{23}\hat{\sigma}_{13}]\mathbf{E}, \\ \partial_t \hat{\sigma}_{13} &= i\omega_{31}\hat{\sigma}_{13} - i[\mathbf{d}_{23}^*\hat{\sigma}_{12} - \mathbf{d}_{12}^*\hat{\sigma}_{23}]\mathbf{E}, \\ \partial_t \hat{\sigma}_{23} &= i\omega_{32}\hat{\sigma}_{23} - i[\mathbf{d}_{23}^*(\hat{\sigma}_{22} - \hat{\sigma}_{33}) - \mathbf{d}_{12}\hat{\sigma}_{13}]\mathbf{E} \end{aligned} \quad (2.15)$$

for the atomic coherences. $\omega_{jk} = \omega_j - \omega_k$ is the $|j\rangle \rightarrow |k\rangle$ transition frequency.

There is a number of reasonable approximations one can apply to simplify equations 2.14 and 2.15. First of all, I will assume that we are only interested in storing single or few photons, which means I assume no significant changes to the atomic populations (the storage ensemble is composed of $\sim 10^{10}$ atoms, see section 3.1.1), i.e. $\hat{\sigma}_{11} = 1$ and $\hat{\sigma}_{22} = \hat{\sigma}_{33} = 0$ ⁴. This allows me to drop equations 2.14 altogether and only focus on the atomic coherences.

Secondly, I will assume the signal $\mathbf{E}_s$ and control $\mathbf{E}_c$ fields only couple to the $|1\rangle \rightarrow |2\rangle$ and $|2\rangle \rightarrow |3\rangle$ transitions, respectively. In the ORCA memory this is completely reasonable due to the large frequency difference between the two transitions. In the case of Λ -based

³Since $\sum_{jj} \hat{\sigma}_{jj} = I$, the identity, $\partial_t \hat{\sigma}_{22} = -\partial_t(\hat{\sigma}_{11} + \hat{\sigma}_{33})$.

⁴The initial, thermal occupancy of states $|2\rangle$ and $|3\rangle$ is negligible, due to the cascaded structure of the optical, high energy transitions.

protocols, coupling of the control $\mathbf{E}_c$ to the $|1\rangle \rightarrow |2\rangle$ transition can lead to four-wave mixing (FWM) noise and gain [139]. This outlines one of the major benefits of ORCA, in that it is immune to FWM (see discussion in section 1.4).

The third approximation I will apply to equations 2.15 is the rotating wave approximation (RWA) [23]. I start by transforming $\hat{\sigma}_{jk}$ to the rotating frame (dropping the operator symbol):

$$\tilde{\sigma}_{jk} = \hat{\sigma}_{jk} e^{i\omega_{jk}(t-z/c)}. \quad (2.16)$$

Equations 2.15 then transform to:

$$\begin{aligned} \partial_t \tilde{\sigma}_{12} &= -i(\mathbf{d}_{12}^* e^{-i\omega_{21}(t-z/c)} + \mathbf{d}_{23} \tilde{\sigma}_{13} e^{i\omega_{23}(t-z/c)}) \mathbf{E}, \\ \partial_t \tilde{\sigma}_{13} &= -i(\mathbf{d}_{23}^* \tilde{\sigma}_{12} e^{-i\omega_{23}(t-z/c)} - \mathbf{d}_{12}^* \tilde{\sigma}_{23} e^{-i\omega_{21}(t-z/c)}) \mathbf{E}, \\ \partial_t \tilde{\sigma}_{23} &= i\mathbf{d}_{12} \tilde{\sigma}_{13} e^{i\omega_{21}(t-z/c)} \mathbf{E}. \end{aligned} \quad (2.17)$$

Now, as I insert the full expression for $\mathbf{E}$, we will notice that I will get terms oscillating at the sum of optical frequencies, and at a frequency equal to the detuning Δ of the fields to their respective transitions (we assume both fields are detuned by the same amount, i.e. $\omega_{21} - \omega_s = \omega_{32} - \omega_c = \Delta$). For example, for $\tilde{\sigma}_{23}$:

$$\begin{aligned} \partial_t \tilde{\sigma}_{23} &= i\mathbf{d}_{12} \tilde{\sigma}_{13} e^{i\omega_{21}(t-z/c)} \left(i\mathbf{v}_s g_s S e^{i\omega_s(t-z/c)} - i\mathbf{v}_s^* g_s S^\dagger e^{-i\omega_s(t-z/c)} \right. \\ &\quad \left. + \mathbf{v}_c E_c e^{i\omega_c(t+z/c)} + \mathbf{v}_c^* E_c^* e^{-i\omega_c(t+z/c)} \right) \\ &= -\mathbf{d}_{12} \mathbf{v}_s g_s S \tilde{\sigma}_{13} e^{i(\omega_{21} + \omega_s)(t-z/c)} + \mathbf{d}_{12} \mathbf{v}_s^* g_s S^\dagger \tilde{\sigma}_{13} e^{i(\omega_{21} - \omega_s)(t-z/c)} \\ &\quad + i\mathbf{d}_{12} \mathbf{v}_c E_c \tilde{\sigma}_{13} e^{i\omega_{21}(t-z/c) + i\omega_c(t+z/c)} + i\mathbf{d}_{12} \mathbf{v}_c^* E_c^* \tilde{\sigma}_{13} e^{i\omega_{21}(t-z/c) - i\omega_c(t+z/c)} \\ &\approx \mathbf{d}_{12} \mathbf{v}_s g_s S^\dagger \tilde{\sigma}_{13} e^{i\Delta(t-z/c)} \end{aligned} \quad (2.18)$$

As one can see, I have neglected the fast oscillating terms, and terms arising from cross-coupling of the fields to the other transitions, i.e. involving for example $\mathbf{d}_{12} \mathbf{v}_c E_c$. This is the essence of the RWA. Doing the same for the rest of equations 2.17, I end up with:

$$\begin{aligned}
\partial_t \tilde{\sigma}_{12} &\approx \mathbf{d}_{12}^* \mathbf{v}_s g_s S e^{-i\Delta(t-z/c)} - i \mathbf{d}_{23} \mathbf{v}_c \tilde{\sigma}_{13} E_c e^{-i\Delta(t-z/c)} e^{-2ik_c z}, \\
\partial_t \tilde{\sigma}_{13} &\approx -i \mathbf{d}_{23}^* \mathbf{v}_c^* \tilde{\sigma}_{12} E_c^* e^{i\Delta(t-z/c)} e^{2ik_c z} - \mathbf{d}_{12}^* \mathbf{v}_s \tilde{\sigma}_{23} g_s S e^{-i\Delta(t-z/c)}, \\
\partial_t \tilde{\sigma}_{23} &\approx \mathbf{d}_{12} \mathbf{v}_s g_s S^\dagger \tilde{\sigma}_{13} e^{i\Delta(t-z/c)},
\end{aligned} \tag{2.19}$$

where $2k_c = 2\omega_c/c$ is simply the longitudinal momentum of the control field $\mathbf{E}_c$ in the reference frame of the moving signal field. Before I move on, I will perform one last simplification of equations 2.19, namely the linear approximation [23]. Again, I assume the signal field is composed of only a few photons, meaning that the coherences produced by them will be small. I can thus assume that any terms that include the product of the two will be much smaller than the rest, and thus can be neglected. I then arrive at:

$$\begin{aligned}
\partial_t \tilde{\sigma}_{12} &\approx \mathbf{d}_{12}^* \mathbf{v}_s g_s S e^{-i\Delta(t-z/c)} - i \mathbf{d}_{23} \mathbf{v}_c \tilde{\sigma}_{13} E_c e^{-i\Delta(t-z/c)} e^{-2ik_c z}, \\
\partial_t \tilde{\sigma}_{13} &\approx -i \mathbf{d}_{23}^* \mathbf{v}_c^* \tilde{\sigma}_{12} E_c^* e^{i\Delta(t-z/c)} e^{2ik_c z}.
\end{aligned} \tag{2.20}$$

Note how the linear approximation eliminated the equation for the $\tilde{\sigma}_{23}$ coherence altogether.

In order to remove any remaining oscillating phase factors from the equations, I use the trick of transforming $\tilde{\sigma}_{jk} = e^{i(\theta_j - \theta_k)} \sigma_{jk}$. Now I can insert them in the equations and choose θ_j, θ_k such that they cancel with the phase factors I wish to remove. Using this trick, I find the following transformations to be optimal:

$$\begin{aligned}
\tilde{\sigma}_{12} &= \sigma_{12} e^{-i\Delta(t-z/c)}, \\
\tilde{\sigma}_{13} &= \sigma_{13} e^{2ik_c z},
\end{aligned} \tag{2.21}$$

thus obtaining the following set of equations for the coherences:

$$\begin{aligned}
\partial_t \sigma_{12} &= i\Delta \sigma_{12} + \mathbf{d}_{12}^* \mathbf{v}_s g_s S - i \mathbf{d}_{23} \mathbf{v}_c \sigma_{13} E_c, \\
\partial_t \sigma_{13} &= -i \mathbf{d}_{23}^* \mathbf{v}_c^* \sigma_{12} E_c^*.
\end{aligned} \tag{2.22}$$

These are known as the optical Bloch equations and they describe the effect the optical fields have on the atomic coherences. They do not yet include dephasing of the atomic coherences themselves due to spontaneous emission or the Doppler effect (which will be

the major source of dephasing for the ORCA memory, see section 3.2). I incorporate the former by inserting dephasing terms into the equations, where γ_{21} and γ_{32} are the inverse of the spontaneous emission lifetimes of states $|2\rangle$ and $|3\rangle$, respectively:

$$\begin{aligned}\partial_t \sigma_{12} &= -\frac{\gamma_{21}}{2} \sigma_{12} + i\Delta \sigma_{12} + \mathbf{d}_{12}^* \mathbf{v}_s g_s S - i\mathbf{d}_{23} \mathbf{v}_c \sigma_{13} E_c, \\ \partial_t \sigma_{13} &= -\frac{\gamma_{32}}{2} \sigma_{13} - i\mathbf{d}_{23}^* \mathbf{v}_c^* \sigma_{12} E_c^*.\end{aligned}\tag{2.23}$$

This approach is in accordance with implementing Markovian decay in the Heisenberg picture using Langevin equations and ignoring the Langevin noise operators, which would describe independent fluctuations of the coherences. These fluctuations average to zero, when we consider only the memory efficiencies [23], so it is safe to ignore them for our purposes.

2.1.5 Maxwell's equations

While the previous section derives the equations describing the behaviour of the atomic coherences due to interaction with the signal and control fields, for memory applications one is more interested in how the atoms influence the signal field, which we wish to store. The classic way of describing how light propagates through a medium is via Maxwell's equations (with no free charge or current):

$$\begin{aligned}\nabla \cdot \mathbf{D} &= 0, \quad \nabla \cdot \mathbf{B} = 0, \\ \nabla \times \mathbf{E} &= -\partial_t \mathbf{B}, \quad \nabla \times \mathbf{H} = \partial_t \mathbf{D},\end{aligned}\tag{2.24}$$

where the magnetic field $\mathbf{B} = \mu_0 \mathbf{H}$, $\mathbf{D} = \epsilon_0 \mathbf{E} + \mathbf{P}$, $\mathbf{P}$ being the atomic polarisation density, and $\epsilon_0 \mu_0 = 1/c$, c being the speed of light in vacuum [23]. Using the identity $\nabla \times (\nabla \times \mathbf{E}) = \nabla \cdot (\nabla \cdot \mathbf{E}) - \nabla^2 \cdot \mathbf{E}$ and ignoring the propagation dynamics of the control field $\mathbf{E}_c$ ⁵, it is straightforward to derive the wave equation for the signal field $\mathbf{E}_s$:

$$\left[\nabla^2 - \frac{1}{c^2} \partial_t^2 \right] \mathbf{E}_s = \mu_0 \partial_t^2 \mathbf{P}_s,\tag{2.25}$$

⁵I assume the strong control field only couples to the empty transition, and thus its propagation is not significantly affected by the atomic medium.

where $\mathbf{P}_s$ is the component of the atomic polarization which acts as a source for the signal field, that can be re-written in terms of the slowly varying atomic polarisation $\mathbf{P}_s = \tilde{\mathbf{P}}_s e^{i\omega_s(t-z/c)}$. Inserting the expression for the signal field (equation 2.6) into equation 2.25 and applying the slowly varying envelope approximation⁶ then gives [23]:

$$\left(\partial_z + \frac{1}{c}\partial_t\right)S = -\frac{\mu_0\omega_s^2}{2g_s k_s} \mathbf{v}_s^* \cdot \tilde{\mathbf{P}}_s, \quad (2.26)$$

where I have also dropped the term involving the transverse Laplacian $\nabla_\perp^2 = \partial_x^2 + \partial_y^2$, since we are only interested in propagation in one dimension with no diffraction (paraxial approximation [23]).

2.1.6 Maxwell-Bloch equations

Before I combine equations 2.23 and 2.26, I will define the continuous collective polarisation operator P and the orbital wave operator B by summing the contributions to the coherences σ_{jk} across the atomic ensemble:

$$\begin{aligned} P &= \frac{1}{\sqrt{nA\delta z}} \sum_{\beta(z)} \sigma_{12}^\beta, \\ B &= \frac{1}{\sqrt{nA\delta z}} \sum_{\beta(z)} \sigma_{13}^\beta, \end{aligned} \quad (2.27)$$

where n is the number density of the atomic ensemble and the index β runs over all atoms within the small ensemble slice of volume $A\delta z$ at position z (I treat the atomic medium as one dimensional), and A is the beam area. One can see, as $A\delta z \rightarrow 0$, P and B will simply describe single atom coherences. On the other hand, according to the Holstein-Primakhoff approximation, we can treat the collective polarisation and orbital wave as bosonic modes, carrying numerable atomic excitations [153].

The combined set of Maxwell-Bloch equations describing the light-matter interaction of the memory is thus:

⁶E.g. $\left|k_s \left(\partial_z + \frac{1}{c}\partial_t\right)S\right| \gg \left|\left(\partial_z^2 + \frac{1}{c^2}\partial_t^2\right)S\right|$.

$$\begin{aligned}
\partial_t S &= -c\partial_z S - c\kappa^* P, \\
\partial_t P &= -(\frac{\gamma_{21}}{2} - i\Delta)P + \kappa S - i\Omega B, \\
\partial_t B &= -\frac{\gamma_{32}}{2}B - i\Omega^* P,
\end{aligned} \tag{2.28}$$

where the coupling constant κ is given by:

$$\kappa = \frac{\mathbf{d}_{12}^* \cdot \mathbf{v}_s}{\hbar} g_s \sqrt{n} = \frac{\mathbf{d}_{12}^* \cdot \mathbf{v}_s}{\hbar} \sqrt{\frac{\hbar \omega_s n}{2\epsilon_0 c}} \tag{2.29}$$

and the Rabi frequency of the control field is:

$$\Omega = \frac{\mathbf{d}_{23}^* \cdot \mathbf{v}_c}{\hbar} E_c. \tag{2.30}$$

We can see that equations 2.28 represent a transformation between a travelling optical bosonic mode S and a stationary atomic bosonic mode B (via a polarisation P). These equations thus fully characterise the memory interaction, as outlined in the beginning of this chapter, and describe the process of an optical signal field S being stored as an orbital wave B , mediated by a control field represented by Ω . It is important to realise here that, because of the counter-propagating nature of the control field in ORCA, Ω is in fact strongly dependent on time and space, i.e. $\Omega = \Omega(t + z/c)$. This prevents me from straightforwardly transforming equations 2.28 into the reference frame moving with the signal pulse.

2.2 Analytic solution

Because the signal and control fields are counter-propagating in the ORCA memory, and thus there exists no obvious frame of reference where they are both stationary, it is not straightforward to solve equations 2.28 analytically. Before attempting this, I thus must make some further approximations. I will start by adiabatically eliminating the excited state P , setting $\partial_t P = 0$. This is reasonable when the natural dynamics of the optical polarization P are overwhelmed by the motion driven by the signal and control fields, also known as adiabatic following [23]. I assume here our detuning Δ is much larger than the control Rabi frequency Ω and pulse bandwidth δ . Equations 2.28 then become:

$$\begin{aligned}
(\partial_z + \frac{1}{c}\partial_t)S(z, t) &= \frac{i\kappa^*}{\Gamma}\Omega(t + z/c)B(z, t) - \frac{|\kappa|^2}{\Gamma}S(z, t), \\
\partial_t B(z, t) &= -\frac{i\kappa}{\Gamma}\Omega^*(t + z/c)S(z, t) - \frac{|\Omega(t+z/c)|^2}{\Gamma}B(z, t),
\end{aligned} \tag{2.31}$$

where $\Gamma = \gamma - i\Delta$ (and $\gamma = \frac{\gamma_{21}}{2}$) and I also assumed the orbital wave does not dephase during the interaction, i.e. $\gamma_{32} = 0$. The term $\frac{|\Omega|^2}{\Gamma}$ is responsible for the dynamic AC Stark effect [23]. Note that I have explicitly written down the time and space dependence of S, B, Ω , with $\Omega(t + z/c)$ accounting for the temporal profile of a counter-propagating control field.

2.2.1 Short cell limit

I will start by considering the case when the length of the control pulse is much larger than the length of the atomic medium. If I transform into the frame moving with the signal pulse, i.e. $x = t - z/c$, we notice that the control field will have the form $\Omega(x + 2z/c)$, where $2z/c$ is the signal/control pulse walk-off. The maximum walk-off in a vapour cell of length L can therefore be $2L/c$. If the control pulse length ($c\tau_c$, where τ_c is the pulse duration) is now much longer than L , i.e. the control pulse shape does not change significantly over a distance L , I can approximate:

$$\Omega(x + 2z/c) \approx \Omega(x + 2L/c) \approx \Omega(x) \tag{2.32}$$

Equations 2.31 then take the form (in the frame of the moving signal pulse⁷):

$$\begin{aligned}
\partial_z S(z, x) &= \frac{i\kappa^*}{\Gamma}\Omega(x)B(z, x) - \frac{|\kappa|^2}{\Gamma}S(z, x), \\
\partial_x B(z, x) &= -\frac{i\kappa}{\Gamma}\Omega^*(x)S(z, x) - \frac{|\Omega(x)|^2}{\Gamma}B(z, x),
\end{aligned} \tag{2.33}$$

These are precisely the equations describing the co-propagating Raman memory interaction (in the adiabatic approximation). Thankfully, these have been solved analytically before in normalised units⁸ [23] and so I will only summarise the derivation of the solution here. I make the following transformation of variables:

⁷Notice how $\partial_z|_t + \frac{1}{c}\partial_t|_z = \partial_z|_x$ and $\partial_t|_z = \partial_x|_z$.
⁸ $S \rightarrow \frac{S}{\sqrt{\gamma}}$ and $B \rightarrow \frac{B}{\sqrt{L}}$.

$$\begin{aligned}
\alpha(z, m)e^{-(Mm+dz)/\Gamma} &= \sqrt{M} \frac{S(z, x)}{\Omega(x)}, \\
\beta(z, m)e^{-(Mm+dz)/\Gamma} &= B(z, x).
\end{aligned}
\tag{2.34}$$

m here is the normalised integrated Rabi frequency $m(x) = \frac{1}{M} \int_{-\infty}^x dx |\Omega(x)|^2$, the normalisation given by $M = \int_{-\infty}^{\infty} dx |\Omega(x)|^2$. Equations 2.34 then take the form:

$$\begin{aligned}
\partial_z \alpha(z, m) &= -C \beta(z, m), \\
\partial_m \beta(z, m) &= C \alpha(z, m),
\end{aligned}
\tag{2.35}$$

where the Raman coupling constant C is given by

$$C = \frac{\sqrt{M}|\kappa|}{\Gamma} \approx \frac{\sqrt{M}|\kappa|}{\Delta} \approx \frac{\Omega_0 \sqrt{\tau_c} |\kappa|}{\Delta} \propto \frac{\Omega_0}{\sqrt{\delta}} \frac{\sqrt{n}}{\Delta} \tag{2.36}$$

in the limit $\Delta \gg \gamma$. Ω_0 is the peak Rabi frequency of the control pulse. Note how the coupling constant C depends on the atomic number density n , which shows how the memory interaction is collectively enhanced by using an ensemble of atoms.

Equations 2.35 describe a beamsplitter-like interaction between a travelling optical mode and a stationary atomic mode, where the reflectivity/transmittivity of the beam-splitter can be dynamically controlled by manipulating the Raman coupling constant C . The solution to Eq. 2.35 will have the form:

$$\begin{aligned}
\alpha(z, m) &= \alpha_0(m) - C \int_0^m \sqrt{\frac{z}{m-m'}} J_1[2C\sqrt{z(m-m')}] \alpha_0(m') dm' \\
&\quad - C \int_0^z J_0[2C\sqrt{(z-z')m}] \beta_0(z') dz', \\
\beta(z, m) &= \beta_0(z) - C \int_0^z \sqrt{\frac{m}{z-z'}} J_1[2C\sqrt{(z-z')m}] \beta_0(z') dz' \\
&\quad + C \int_0^m J_0[2C\sqrt{z(m-m')}] \alpha_0(m') dm',
\end{aligned}
\tag{2.37}$$

J_0 and J_1 being the zeroth and first order Bessel functions respectively [23]. The solution can also be written in the form:

$$\begin{aligned}
S_{out}^j &= \mu_j S_{in}^j + \lambda_j B_{in}^j, \\
B_{out}^j &= \mu_j B_{in}^j + \lambda_j S_{in}^j.
\end{aligned}
\tag{2.38}$$

The index j describes the mode-number of the optical signal field S and orbital wave B . The constants μ_j and λ_j , which couple these particular optical and atomic modes can be found by performing a singular value decomposition (SVD) of the transformation described by Eq. 2.37, and, squared, will correspond to the storage/retrieval efficiencies for those particular modes [23]. It has been shown that for small C the Raman memory is predominantly single-mode [136] and the storage efficiency scales as C^2 [134]. In the regime where the length of the control pulse is much larger than the length of the atomic medium, the ORCA memory should therefore exhibit similar behaviour.

In terms of the achievable memory efficiency, it has been shown that with appropriate mode-matching (via e.g. pulse shaping of the signal field) the co-propagating Raman memory should be able to reach unit efficiencies [154]. Unfortunately, that analysis did not include the experimentally unavoidable coupling of the control field to the populated transition in the Raman case, which leads to gain on the signal at high coupling strengths, because of FWM [100]. This precludes the Raman memory reaching high efficiencies with low noise⁹. As FWM is altogether absent in the ORCA protocol, in the short cell regime, where ORCA behaves analogously to the Raman memory (but with no FWM), it should therefore actually be possible to reach unit efficiencies without deteriorating the noise properties.

2.2.2 Long cell limit

The opposite limit one can investigate is when the optical pulses are much shorter than the storage medium. In order to solve equations 2.31 in the limit of a long storage medium, i.e. in which the pulses are able to completely pass each other in the atomic ensemble, I will assume negligible linear dispersion and absorption of our signal field (reasonable when $\Delta \gg \gamma$), and neglect the AC Stark effect due to the control field, which means $\frac{|\Omega|^2}{\Delta} \ll (\gamma, \delta) \ll \Delta$, which is consistent with the adiabatic approximation I have made before. We are then left with solving:

$$\begin{aligned} (\partial_z + \frac{1}{c}\partial_t)S(z, t) &= \frac{i\kappa^*}{\Gamma}\Omega(t + z/c)B, \\ \partial_t B(z, t) &= -\frac{i\kappa}{\Gamma}\Omega^*(t + z/c)S. \end{aligned} \tag{2.39}$$

⁹Unless methods to suppress FWM are employed [140, 141].

To do so, I start with a coordinate transformation $x = t - z/c$ and $y = t + z/c$, and define $\alpha(x, y) = S\left(\frac{(y-x)c}{2}, \frac{(y+x)}{2}\right)$ and $\beta(x, y) = B\left(\frac{(y-x)c}{2}, \frac{(y+x)}{2}\right)$, obtaining:

$$\begin{aligned}\partial_y \alpha(x, y) &= \frac{ic\kappa^*}{2\Gamma} \Omega(y) \beta(x, y), \\ (\partial_x + \partial_y) \beta(x, y) &= -\frac{i\kappa}{\Gamma} \Omega^*(y) \alpha(x, y).\end{aligned}\tag{2.40}$$

In order to get rid of the ∂_x derivative, I Fourier transform eq. 2.40, defined as $\tilde{f}(k, y) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(x, y) e^{ikx} dx$. I get:

$$\begin{aligned}\partial_y \tilde{\alpha}(k, y) &= \frac{ic\kappa^*}{2\Gamma} \Omega(y) \tilde{\beta}(k, y), \\ \partial_y \tilde{\beta}(k, y) &= ik \tilde{\beta}(k, y) - \frac{i\kappa}{\Gamma} \Omega^*(y) \tilde{\alpha}(k, y).\end{aligned}\tag{2.41}$$

I now define a new coordinate:

$$w(y) = \frac{1}{W} \int_{-\infty}^y \Omega(y') dy',\tag{2.42}$$

where $W = \int_{-\infty}^{\infty} \Omega(y) dy$ is the area of the control pulse. $w \in [0, 1]$ can be interpreted as a clock for the interaction between the control and signal fields, or the extent to which the control pulse has passed the medium. In the new coordinate system, the equations of motion become¹⁰:

$$\begin{aligned}\partial_w \tilde{\alpha}(k, w) &= \frac{icW\kappa^*}{2\Gamma} \tilde{\beta}(k, w), \\ \partial_w \tilde{\beta}(k, w) &= \frac{iWk}{\Omega(y)} \tilde{\beta}(k, w) - \frac{iW\kappa}{\Gamma} \tilde{\alpha}(k, w),\end{aligned}\tag{2.43}$$

where I additionally assumed $\Omega(y) = \Omega^*(y)$. Redefining $\tilde{\alpha}(k, w) = \tilde{\alpha}_0(k, w)$ and $\tilde{\beta}(k, w) = \tilde{\beta}_0(k, w) e^{iky}$, I get:

$$\begin{aligned}\partial_w \tilde{\alpha}_0(k, w) &= \frac{icW\kappa^*}{2\Gamma} \tilde{\beta}_0(k, w) e^{iky}, \\ \partial_w \tilde{\beta}_0(k, w) &= -\frac{iW\kappa}{\Gamma} \tilde{\alpha}_0(k, w) e^{-iky}.\end{aligned}\tag{2.44}$$

After an inverse Fourier transform:

$$\begin{aligned}\partial_w \alpha_0(x, w) &= \frac{icW\kappa^*}{2\Gamma} \beta_0(x - y, w), \\ \partial_w \beta_0(x, w) &= -\frac{iW\kappa}{\Gamma} \alpha_0(x + y, w).\end{aligned}\tag{2.45}$$

¹⁰Utilising the chain rule $\partial_y = (\partial_y w) \partial_w$.

The extra phase factor in Fourier space is equivalent to a shift in space. Hence, $\alpha(x, w) = \alpha_0(x, w)$ and $\beta(x, w) = \beta_0(x - y, w)$. The equations of motion then take the symmetric form:

$$\begin{aligned}\partial_w \alpha(x, w) &= \frac{icW\kappa^*}{2\Gamma} \beta(x, w), \\ \partial_w \beta(x, w) &= -\frac{iW\kappa}{\Gamma} \alpha(x, w).\end{aligned}\tag{2.46}$$

By differentiating both sides of each equation and inserting the non-differentiated form, Eq. 2.49 can be re-written in the easy-to-solve form:

$$\begin{aligned}\partial_w^2 \alpha(x, w) &= C'^2 \alpha(x, w), \\ \partial_w^2 \beta(x, w) &= C'^2 \beta(x, w).\end{aligned}\tag{2.47}$$

C' being the ORCA coupling parameter, which characterises the interaction strength between modes α and β :

$$C' = \sqrt{\frac{c}{2}} \frac{W|\kappa|}{\Gamma} \approx \sqrt{\frac{c}{2}} \frac{W|\kappa|}{\Delta} \propto \frac{\Omega_0}{\delta} \frac{\sqrt{n}}{\Delta}\tag{2.48}$$

in the limit $\Delta \gg \gamma$. Solving equations 2.49, I obtain:

$$\begin{aligned}\alpha(x, w) &= \alpha_{in}(x, w)|_{w=0} \cos(C'w) + \beta_{in}(x, w)|_{w=0} \sin(C'w), \\ \beta(x, w) &= \beta_{in}(x, w)|_{w=0} \cos(C'w) + \alpha_{in}(x, w)|_{w=0} \sin(C'w).\end{aligned}\tag{2.49}$$

For example, in the storage process, $\beta_{in} = 0$ and so $\beta_{out}|_{w=1} = \alpha_{in}|_{w=0} \sin C'$. $w = 0$ signifies the state of the system before the interaction, i.e. $y = t + z/c \rightarrow -\infty$, and similarly $w = 1$ is after the signal and control pulses interacted, i.e. $y = t + z/c \rightarrow \infty$. If I define the input signal $S_{in}(t)$ at $z \rightarrow -\infty$, and the output orbital wave $B_{out}(z)$ for $t \rightarrow \infty$. Then:

$$\left| \int_{-\infty}^{\infty} B_{out}(z) dz \right|^2 = \sin^2 C' \left| \int_{-\infty}^{\infty} S_{in}(t) dt \right|^2\tag{2.50}$$

and so the storage efficiency is $\eta_{in} = \sin^2 C'$. The solution 2.50 could have been expected from looking at equations 2.49, which describe a light-matter beamsplitter-like interaction, similarly to the short cell case (though defined for different variables). The retrieval

efficiency will also be $\eta_{out} = \sin^2 C'$, and thus the total memory efficiency scales as $\eta_{tot} = \sin^4 C'$. This is interesting for two reasons. Firstly, for small C' the storage/retrieval efficiencies in the long cell regime therefore scale as $C'^2 \propto n\left(\frac{\Omega_0}{\Delta}\right)^2$, similarly to the case of the short cell discussed in section 2.2.1, where they scaled as $C^2 \propto n\left(\frac{\Omega_0}{\Delta}\right)^2$. For larger C and C' , the memory efficiency starts oscillating in the long cell regime, whereas it would saturate in the short cell case. While unsurprising for a beamsplitter-type interaction, physically the oscillation in efficiency in the long-cell limit can be understood as the control field immediately retrieving the only-stored signal, as we overdrive the memory interaction, or in other words, the tail of the control pulse retrieving what the front stored¹¹.

Secondly, according to the solution, the ORCA memory in the long cell regime should in principle be able to reach 100% efficiencies. Of course, at the high coupling strengths required for reaching $\sin C' \sim 1$ the currently ignored effects of linear absorption (and re-absorption) and the AC Stark effect start playing a role, reducing the achievable efficiencies (as will be shown in section 2.3.1).

Another interesting observation is that in the long cell regime, when the signal and control pulses pass each other fully, the control field completely maps the shape of the signal field into the spatially distributed orbital wave¹². This suggests that, in principle, the memory is temporally multimode in this regime, and thus allows the storage of complex pulse shapes, or multiple pulses at different positions in the storage medium.

I have thus found analytical solutions to the equations governing the ORCA memory (Eq. 2.28) in two regimes. In section 2.2.1, I investigated the short cell limit, i.e. the regime where the storage medium, or interaction region, is short compared to the optical pulse length. In this regime the ORCA memory can be reduced to the co-propagating Raman case (albeit without four-wave mixing) [134, 23], inheriting the latter's achievable unit efficiency, and the modal capacity depending on the coupling strength. The second, long cell limit (section 2.2.2), in which the counter-propagating signal and control pulses are much shorter than the storage medium and thus are able to completely pass each

¹¹This will be better seen in section 2.3.

¹²More technically, performing an SVD on the interaction described by Eq. 2.49 would yield a similar set of equations to Eq. 2.38 for the mode transformation. The matrix form of that transformation would simply be the identity matrix multiplied by $\sin C'$, i.e. all modes having the same storage efficiency.

other within the latter, required a new derivation. Neglecting linear absorption and the AC Stark effect, I showed that in this regime the control is able to fully map the signal into the storage medium, making the memory highly multimode. The memory efficiency in this case oscillates with the coupling strength.

It would thus be interesting to investigate the intermediate cell length regime, without neglecting linear absorption and the AC Stark effect (which is believed to reduce the memory efficiency in practical implementations [17, 140]). Solving Eq. 2.28 is difficult analytically. However, it is relatively straightforward to do it numerically, which is the focus of the next part of this chapter.

2.3 Numerical solution

I use two “families” of codes to solve the equations of motion describing the ORCA memory interaction numerically. First, in subsection 2.3.1 I use a Matlab script, which solves the adiabatically reduced Eq. 2.31 in the frame moving with the signal pulse (see section 2.3.2), using a fourth-order Runge-Kutta (RK4) method in the time domain, interspersed with steps in space using Euler’s method (equivalent to RK1) [155].

Generating the results in section 2.3.2, based on large numbers of memory simulations, requires a faster and more accurate method to solve the ORCA equations of motion, and so I use a second, parallel processing program written in Python for this purpose (described in more detail in section 2.3.2). I will assume that the only dephasing experienced by the memory is due to spontaneous emission from the excited state¹³, as described by equations 2.28, which is negligible on the timescales of storage ($\sim$ ns) that I will simulate.

2.3.1 Comparison with analytics

Before embarking on a thorough numerical investigation of the performance of the ORCA memory, we can do a preliminary comparison between the predictions of the analytical solution in section 2.2.2 and a numerical model based on fewer assumptions. This is useful to identify the regime of ORCA operation where the analytical assumptions break down,

¹³Recently, DPhil student Oscar Lazo-Arjona has written simulation code that includes motion-induced dephasing and the effect of a complex hyperfine structure of the storage state in the ORCA memory [146].

and a full numerical model is necessary.

Fig. 2.2 shows a comparison of the results of the analytical treatment in the long cell limit (section 2.2.2) and a numerical model solving Eq. 2.31 for the storage process with, and without the AC Stark effect and linear absorption/dispersion¹⁴. The parameters of the numerical model were chosen such as to closely match the assumptions of the long cell limit, such as large single-photon detuning of the signal/control fields (to remain in the adiabatic condition), as well as long cell length combined with a temporally short control field, so that the latter completely scans the signal pulse inside the memory medium.




Figure 2.2: Comparison of the results of solving Eq. 2.31 numerically with, and without the AC Stark effect and linear absorption/dispersion, and the analytical solution given by Eq. 2.50. Storage efficiency is plotted as a function of the coupling constant C' (as defined in Eq. 2.48).

As we can see, the analytical predictions closely match the more general numerical model without the AC Stark effect and linear absorption (as assumed in the analytical solution in the long cell limit), verifying my derivation in section 2.2.2. For low coupling

¹⁴The storage efficiency in the numerical simulation was found by calculating the ratio of the areas under the simulated signal curve before and after the memory interaction (see section 2.3.2).

$C' < 0.5$, the effect of including the AC Stark effect is barely noticeable, with all three models exhibiting the same behaviour. As C' becomes larger¹⁵, the curve generated by the numerical model with AC Stark starts to diverge considerably from the other two cases shown in Fig. 2.2¹⁶. The oscillation of the memory efficiency with C' becomes faster and the maximum reachable efficiency drops to about half its maximal value without the AC Stark effect. The drop in storage efficiency can be attributed to the AC Stark effect disrupting the two-photon resonance condition (reducing the frequency mode-matching of the signal and control fields), as the strong time-dependent control field shifts the energies of the intermediate and storage states [140].

However, the retrieval process is less sensitive to the AC Stark shift than storage (as we will see in the next section). Thus even in the limit where the memory is overdriven, which I postulate to be the physical process behind the oscillation in efficiency with coupling strength, the retrieval process is relatively easier. This causes the oscillation in memory efficiency to become faster when the AC Stark effect is included, as the weakened storage can no longer compensate for the retrieval as well.

2.3.2 Numerical analysis

In order to gain deeper insight into the behaviour of the ORCA memory in the case of realistic experimental parameters, I use a faster and more accurate code with fewer assumptions to simulate the memory [156]. I model the $6S_{1/2} \rightarrow 6P_{3/2} \rightarrow 6D_{5/2}$ cascade in Cs (which was used for the experiments described in chapters 3 and 4), treating the atomic ensemble as a pure three-level system¹⁷ (ignoring hyperfine splitting). The code simulates the interaction of two counter-propagating Gaussian-shaped electric fields with two atomic coherences in a one dimensional cell of 72 mm length filled with atoms, as described by equations 2.28 (both fields and coherences are complex numbers). I solve the equations in the time domain using a fourth-order Runge-Kutta (RK4) method, calculating the spatial derivative of the signal field $\partial_z S$ using the Chebyshev matrix

¹⁵For the purpose of the simulation I have actually varied the control pulse peak Rabi frequency to change C' , which will induce the strongest AC Stark shift.

¹⁶The effect of linear absorption itself is less significant and mainly causes the maximum achievable storage efficiency to drop by a few %.

¹⁷The values for the dipole moments of the two transitions were taken from Table 3.1, section 3.3.

approach [157].

I choose a cell temperature of 90°C (corresponding to an atomic number density of $n = 8.4 \times 10^{18} \text{ m}^{-3}$), single-photon detuning Δ of 6 GHz and read-in/-out control pulses energy of 70 pJ (pulse bandwidth of 1 GHz, corresponding to a pulse duration of ~ 440 ps). The pulse length in my simulation is on the order of the cell length, i.e. in a regime not covered by the analytical treatment in the previous sections, but similar to the experiments presented in chapters 3 and 4.

Figure 2.3 shows a typical result of simulating both storage and retrieval in the ORCA memory. One can see that when the counter-propagating signal and control fields cross in the memory medium, the signal gets absorbed and an orbital wave (corresponding to a coherence between levels $6S_{1/2}$ and $6D_{5/2}$ or, in other words, states $|1\rangle$ and $|3\rangle$ from section 2.1) is created, as expected. Consequently, when a second control pulse is sent to the memory 3.5 ns later, part of that orbital wave is re-transformed into an optical pulse which leaves the memory. In the simulation showed in Fig. 2.3 the storage efficiency is $\eta_{\text{in}} = 54\%$ and retrieval $\eta_{\text{out}} = 50\%$, yielding a total memory efficiency of $\eta = 27\%$ ¹⁸.

Before I set off on optimizing the ORCA memory efficiency over multiple input parameters using an automatic search algorithm, it is worthwhile to investigate how the memory efficiency depends on the specific memory parameters independently, since one expects the simple dependence on the generalised coupling parameter C' from the previous section to be inadequate to understand the realistic memory operation.

One interesting example, present in off-resonant memory schemes [100], is for instance the counterplay between the single-photon detuning of the optical fields from their respective transitions, and the number density of the atoms forming the memory medium. In order to limit linear absorption of the signal and remain in the adiabatic regime of the memory ($\Delta \gg \Omega$), one requires large detunings. However, this leads to a weaker atom-light coupling strength, which can be compensated by increasing the number of atoms taking part in the interaction. Yet this leads to larger (re-)absorption of the signal. The analytic solution derived in section 2.2.2 does not predict this counterplay, as it explicitly assumes the adiabatic condition and lack of linear absorption. It is therefore interesting

¹⁸I calculate the efficiencies by comparing the simulated input, transmitted and retrieved signal pulse energies.




Figure 2.3: Top left: behaviour of signal field ($|S|^2$ from Eq. 2.28) during ORCA storage and retrieval. The dashed white lines represent the propagation of the control field. Top right: behaviour of orbital wave ($|B|$ from Eq. 2.28) during ORCA storage and retrieval. Bottom: time trace of the signal field at the front and back edges of the memory cell. Visible in orange are both the transmitted (earlier time) and retrieved (later time) fields.

to investigate the memory numerically in the regime, where both effects are significant.

I thus simulate the ORCA memory for different combinations of detunings and atomic number densities, the effects of which are shown in Fig. 2.4. As one can see, increasing the number density leads to stronger absorption of the signal, but at some point does not improve its re-emission, as the signal starts suffering severe incoherent (linear) absorption both during read-in and read-out. Interestingly, we see an oscillation in the efficiencies with increasing number density, also predicted by my analytical treatment (section 2.2.2), but in the general simulated case the oscillations dampen due to linear absorption at high number densities. On the other hand, my analytical treatment predicts a similar oscillation with detuning, while my numerical solution is quite flat. What is more surprising is that actually the memory seems to favour large detunings (with appropriate compensation from higher number densities). This suggests that the AC Stark shift ($\Delta_{AC} \propto \frac{\Omega_0^2}{\Delta}$), ignored in my analytical treatment, may in fact be an important limiting factor to memory efficiency, as was the case for the Raman memory [140]. Especially since it is the read-in efficiency (that relies on signal and control being on two-photon resonance, which

the AC Stark shift disrupts) that improves with detuning, while read-out stays roughly the same.




Figure 2.4: Top/bottom: simulated read-in/-out memory efficiencies as a function of single photon detuning of the optical fields Δ and the atomic number density n of the medium.

Another factor that neither of my analytical treatments include is the effect of the signal/control temporal overlap in the memory cell. By moving the point where the pulses cross (changing e.g. the relative delay between signal and control pulses) in the cell, one should be able to manipulate the memory interaction, potentially also influencing its modal capacity¹⁹. I therefore run the ORCA memory simulations for different relative delays between signal and control (where relative delay $\Delta\tau = 0$ corresponds to pulses meeting in the centre of the cell), while also varying the read-in/-out control pulse energy. The results are shown in Fig. 2.5. Positive delays correspond to the control field arriving later, i.e. the pulses meeting closer to the back edge of the cell.

Unsurprisingly, as the pulses meet at the back edge of the cell, localising the orbital wave there, the read-out becomes more efficient, as the retrieved signal pulse has less of the atomic medium to travel and thus has a lower chance of getting re-absorbed (either via resonant or off-resonant absorption). A similar effect was predicted in the Raman memory,

¹⁹By changing the effective length of the medium in which the signal and control pulses can overlap.




Figure 2.5: Top/bottom: simulated read-in/-out memory efficiencies as a function of read-in/-out control pulse energy $E_{\text{in/out}}$ and relative delay $\Delta\tau$ between signal and control fields.

where at low control field energies maximum memory efficiency could only be reached by backward retrieval, with re-absorption limiting the efficiency otherwise [23, 154]. At high control field energies, backward retrieval becomes less important as the control field is basically able to “push” the coherence out through the medium. We see a similar effect in the ORCA simulation, where at high control field energies it becomes beneficial, in terms of read-out, for the pulses to meet closer to the center of the cell, where pulse overlap is maximised. Unsurprisingly, read-in efficiency is optimised when signal and control meet in the center of the cell, as this maximises their interaction length.

Worth noticing is that the oscillation in storage efficiency with control field energy, as one overdrives the memory interaction, is much weaker than one would expect from the analysis in section 2.3.1 (and rather saturates to a value of roughly 60%). As has been mentioned in the previous sections, the oscillation in memory efficiency is due to concurrent processes of storage and immediate retrieval of the signal as it passes through a control pulse with increasing energy. This is visible especially in the top right panel of Fig. 2.6, which depicts how the orbital wave is generated during storage, and immediately retrieved again as light. This effect can also be seen in the temporal shift and shape

change of the transmitted signal (top left and bottom panels of Fig. 2.6). The oscillation in storage efficiency is less pronounced than expected from the analytical treatment, due to the spatial constraints of the memory cell. By decreasing the time the signal and control pulses are overlapped in the cell, this prevents effective retrieval of the just-stored excitation (which is suggested by the fact the degree of oscillation in Fig. 2.5 varies on where the signal and control pulses meet, and thus the orbital wave is generated).




Figure 2.6: Top left: behaviour of signal field ($|S|^2$ from Eq. 2.28) during ORCA storage and retrieval with a high control pulse energy (0.5 nJ). The dashed white lines represent the propagation of the control field. Top right: behaviour of orbital wave ($|B|$ from Eq. 2.28) during ORCA storage and retrieval. Bottom: time trace of the signal field at the front and back edges of the memory cell. Visible in orange are both the transmitted and retrieved fields.

The retrieval efficiency similarly exhibits weak oscillation with control field energy, though here we see that the efficiency does increase at higher control pulse energies. This is because retrieval is less sensitive to the AC Stark effect (since the retrieved signal will simply match the AC Stark-modified control), and thus the latter does not have such strong efficiency-limiting properties as for storage (though again, the position of maximum pulse overlap will change the oscillation properties somewhat). In summary, from the numerical simulation of the ORCA memory we see that the analytical treatment

from section 2.1 does not contain all the behaviour of the protocol, especially at high coupling strengths. To find the optimal, yet realistic, memory performance it is therefore necessary to use the numerical model.

Numerical optimisation of the ORCA memory

In order to find the maximum achievable memory efficiency, I perform a numerical search over the different ORCA memory parameters, by minimising the function $1 - \eta(\Delta\tau, E_{\text{in}}, E_{\text{out}}, \Delta, n)$ using the Nelder-Mead method [158], included in the used Python distribution. Note that I optimise the efficiency independently over the read-in and read-out control pulse energies, as we can see from Fig. 2.5 that maximum storage and retrieval efficiencies can occur for different control pulse energies. Table 2.3.2 summarises the results of this optimisation²⁰.

Pulse delay $\Delta\tau$	67 ps
Read-in energy E_{in}	47 pJ
Read-out energy E_{out}	943 pJ
Detuning Δ	6.5 GHz
Atomic number density n	$4.9 \times 10^{19} \text{ m}^{-3}$
η_{in}	0.85
η_{out}	0.62

Table 2.1: Memory parameters that optimise memory efficiency, found using an optimisation algorithm based on the Nelder-Mead method.

The maximum efficiency found by the optimisation algorithm (for realistic parameters) was $\eta_{\text{max}} \approx 62\%$ for a cell temperature $\sim 120^\circ\text{C}$. Reaching close to unit efficiencies on the other hand would require optimisation over the pulse shape, which is beyond the scope of the current investigation²¹. Since four-wave mixing gain [100] is absent from ORCA (as

²⁰Worth mentioning is that Δ and n were respectively capped at 12 GHz and $\sim 5 \times 10^{19} \text{ m}^{-3}$, during the optimisation process, and thus higher efficiencies could be reached in principle. Yet for $n > 4.5 \times 10^{19} \text{ m}^{-3}$, self-broadening of the signal transition due to dipole-dipole interactions between atoms [159] becomes comparable to the natural linewidth and should no longer be ignored (as it is in my current analysis).

²¹Memory efficiencies greater than 80% have been seen using a different code, which actually runs the ORCA simulation for different input modes, in order to calculate the Green's function for the ORCA memory interaction.

mentioned in section 2.1.4), it is feasible to reach these high efficiencies while preserving the noise-free properties of the memory.

My numerical analysis of the ORCA memory has shown that it is non-trivial to reach the unit efficiencies predicted by the analytical model, and that in order to simulate the memory properly, one does need to include additional effects such as the AC Stark shift. Nonetheless, even without complicated mode-matching, it should be possible to reach memory efficiencies greater than 50% ²², opening the route to using ORCA for quantum information processing applications (as discussed in section 1.2).

It would therefore be desirable to implement the ORCA memory experimentally, not least to verify its noise-free properties, which my current theoretical model does not include. This is the topic of the next two chapters.

²²It is however important to remember that the optimal efficiency I found corresponds essentially to a zero storage time (on the timescale of the spontaneous emission lifetime), and that in practice, additional dephasing mechanisms will limit the memory efficiency (investigated in section 3.2).

3

Classical characterisation of the ORCA memory in warm caesium vapour

In order to test the predictions of the previous chapter and verify the noise-free nature of ORCA we implement the protocol in warm Cs vapour with GHz-bandwidth near-infrared light [146]. The choice of atomic system and signal wavelength was determined by pre-existing lab infrastructure for the Raman memory in Cs [16, 100] (constructed by Dr. Dylan Saunders and PhD student Sarah Thomas, and adapted for the ORCA experiment by myself and Dr. Patrick Ledingham). In this chapter, I describe the experimental setup for characterising the “classical” performance of the ORCA memory using weak coherent states, after which I present the results of measuring the essential memory parameters: lifetime, efficiency and noise level.

3.1 Experimental setup

The experimental setup for storing weak coherent states in the ORCA memory can be divided into four parts, i.e. the memory cell itself with surrounding optics, the signal and control generation stages, and the detection scheme, all shown in Fig. 3.1, which I will describe in more detail in the following sections. Consequently, I present evidence of successful ORCA operation.




Figure 3.1: Complete experimental setup for storing weak coherent states in ORCA. (a) Memory stage: DM1, DM2 - dichroic mirrors; HWP - half-wave plate; QWP - quarter-wave plate; PBS - polarising beamsplitter; HR@852nm - highly reflective mirror for 852 nm; sPD - slow photodiode. (b) Signal generation stage: ISO - optical isolator; SHG - second-harmonic generation crystal; DM - dichroic mirror; PBD - polarising beam displacer; ND - neutral density filter; FP - Fabry-Pérot etalon; PDC - single-photon generation setup; Exp - memory experiment. (c) Control generation stage: GT - Glan-Thompson prism. (d) Signal detection stage: BP - 10 nm bandpass filter @ 850 nm; APD - avalanche photodiode.

3.1.1 Memory cell

We implement ORCA in warm caesium-133 vapour. The electronic energy structure used is shown in Fig. 3.2. The $6S_{1/2}(F=4)$ level acts as our ground state $|1\rangle$, the $6P_{3/2}(F=3,4,5)$ ¹ manifold as our intermediate $|2\rangle$, and the $6D_{5/2}(F=2,3,4,5,6)$ manifold as our storage state $|3\rangle$. It has been shown that such a complex atomic structure can be reduced to an effective three-level system in the case of broadband excitation² [161]. We chose the D2 line of caesium for our signal wavelength as our single-photon source operated at 852 nm (chapter 4). As for the choice of the control transition, the $6P_{3/2} \rightarrow 6D_{5/2}$ transition lies closer in wavelength to the signal one, compared to the $6P_{3/2} \rightarrow 6D_{3/2}$ [162], which leads to slightly reduced motion-induced dephasing (see section 3.2.1). More importantly though, the $6P_{3/2} \rightarrow 6D_{5/2}$ transition has a three times larger dipole moment than the corresponding $6D_{3/2}$ transition [163].




Figure 3.2: Level scheme in caesium [162] employed in the described ORCA memory experiment, including hyperfine splitting.

The choice of caesium was motivated from the benefits it provided in previous experiments with the Raman memory protocol (see section 1.3.2). These arise from caesium being the heaviest stable alkali metal. First, a larger mass leads to slower diffusion of caesium atoms in the vapour, which was found to be an important limiting factor to Λ - type memory lifetime in warm vapour [100]. Secondly, caesium has a large nuclear

¹The $6S_{1/2} \rightarrow 6P_{3/2}$ transition is known as the caesium D2 line, whereas the $6S_{1/2}(F=3) \rightarrow 6S_{1/2}(F=4)$ is the caesium atomic clock transition [160].

²Though this approximation will not be as accurate when we consider memory lifetime, section 3.2.2.

spin, which leads to a large hyperfine splitting of the ground state of 9.2 GHz. This in turn sets the fundamental limit to the storage bandwidth in Λ -type quantum memories [17]. An additional benefit of Cs, which also applies to ORCA, is its high vapour pressure, making it relatively easy to achieve large number densities of atoms at moderately high temperatures.

The atomic number density can be extracted from the vapour’s optical depth, which we determine by scanning the frequency of a weak continuous wave (CW) beam and measuring its transmission through the memory cell on a slow photodiode (see [Appendix A. Optical depth](#) for more information about measuring the optical depth).

As the storage medium we use an off-the-shelf 71.8 mm -long borosilicate Cs reference cell, which we heat up to roughly 91°C³ (more on measuring the cell temperature in [Appendix A. Optical depth](#)). The temperature was chosen to have a sufficiently high number density ($n \sim 10^{19} \text{ m}^{-3}$) to provide good light-matter coupling, without observing significant self-broadening of the signal transition (see section 2.3.2). To heat the cell up, we use an oven made of three concentric layers of teflon and μ -metal, where the first layer of metal is wrapped in quad-twisted heating wire⁴. Such infrastructure was relevant for suppressing magnetic dephasing in the Raman memory [100], though is less relevant in ORCA.

Since the storage state in ORCA is a doubly-excited electronic state it has no thermal excitations even at high temperatures. This means the ORCA memory does not require optical pumping to empty the storage state, as is the case for most Λ -type memories [100]. Without the need for optical pumping, we could operate the Cs cell at high temperature without a buffer gas⁵. Alkali-buffer atom collisions, while spin-preserving (hence used to extend the lifetime in Λ -type memories⁶), are not orbital-preserving and thus would lead to a deterioration in the ORCA memory lifetime (to on the order of the inverse of the pressure broadening of the atomic transition, i.e. tens of MHz per Torr of buffer gas

³Keeping the cell temperature under 100°C has the added value of extending the vapour cell lifetime, as Cs becomes more reactive to the glass at higher temperatures.

⁴Quad-twisted wire should prevent the production of magnetic fields arising from a current being run through it, which could otherwise contribute to magnetically-induced dephasing of the memory [149].

⁵Optical pumping in Λ -type memories is limited at high temperatures, because of light trapping, which can be somewhat suppressed using a molecular buffer gas [100].

⁶Buffer gas decreases the speed of atomic diffusion in warm vapour, extending the transit time of atoms through the optical beam. The latter limits the memory lifetime in warm vapour memories [100].

[164]).

The cell was surrounded from each side by lenses with 300 mm focal length, focusing both signal and control to $\sim 300 \mu\text{m}$ beam widths roughly in the middle of the cell. The control field (section 3.1.3) was combined with the signal (section 3.1.2) mode, counter-propagating, on a short-pass dichroic mirror, as shown in Fig. 3.1 (a). The spatial overlap of signal and control was optimised daily using two irises before and after the cell. The temporal overlap of the signal and read-in pulses in the Cs cell was controlled using the control laser's lock-to-clock (L2C) electronics (see section 3.1.3).

The signal was prepared in horizontal polarisation before the memory, while the control polarisation was optimised using half- and quarter-wave plates, in order to maximise the measured read-out signal at the start of the experiment (afterwards, it was kept constant throughout the measurements in this chapter)⁷. After passing through the memory, the signal is sent towards the detection stage of the experiment (section 3.1.4). The 917 nm control, on the other hand, is predominantly reflected of a shortpass dichroic mirror towards a beam dump, with a small amount being picked-off at a high-reflectivity mirror at 852 nm. This light is sent onto a fast photodiode, used to monitor the timing of the control field with respect to the signal. The signal transmission from in front of the first memory filter to after the last one was $\sim 40\%$.

3.1.2 Signal

The weak coherent state signal was generated in an actively mode-locked titanium sapphire (Ti:Sapph) laser (Spectra-Physics Tsunami), with an intracavity Gires-Tournois interferometer (GTI), producing pulses of 1.35 GHz bandwidth. The Ti:Sapph was pumped by a 13 W continuous wave laser at 532 nm. The oscillator generated ~ 330 ps long pulses at 80 MHz repetition rate with an average power of 2 W. We could tune the center wavelength of the laser coarsely by rotating a birefringent filter inside the laser cavity. On the other hand, fine tuning of the laser center frequency was achieved by changing the voltage on the GTI (see Appendix B. Laser locking for more details). We used this control to reference-lock the signal frequency 6 GHz to the red of the $6S_{1/2}(F=4) \rightarrow 6P_{3/2}$ transi-

⁷A more detailed discussion of the effects of signal and control polarisations on the ORCA memory can be found in Appendix D. Polarisation.

tion. We chose a detuning of 6 GHz, as it would provide sufficient coupling strength, with negligible ($< 2\%$) linear absorption of the signal. The spectrum of the pulses directly from the laser, measured with a commercial scanning Fabry-Perót interferometer, is shown in Fig. 3.3. The peaks are separated by 80 MHz, i.e. the pulse repetition rate.




Figure 3.3: Signal Ti:Sapph spectrum, measured on a scanning Fabry-Pérot interferometer (SFP). The visibility of the peaks is limited by mode-matching to the SFP.

The signal laser setup is shown in Fig. 3.1 (b). After passing through an optical isolator the signal light is doubled in a bulk periodically-poled potassium titanyl phosphate (ppKTP) crystal. The second harmonic and the unconverted fundamental are then separated on a dichroic mirror with the former sent to the single-photon generation setup (section 4.1.3) and the latter sent into a Pockels cell, triggered by a signal generated from the Ti:Sapph pulse train.

By applying a fast high-voltage pulse to the Pockels cell, we are able to rotate the polarisation of a single Ti:Sapph pulse. This, combined with a polarising beam displacer, allows us to pick a single pulse out of the 80 MHz pulse train at a maximum rate of 1 MHz and extinction ratio of $\sim 20,000:1$. The unpicked pulses are sent into a beam dump, while the picked pulse is passed through a Fabry-Pérot etalon to reduce its bandwidth to 0.81 GHz, and a series of neutral density filters to reduce its amplitude down to the

single-photon level. After this treatment, the weak coherent state signal is coupled to single-mode fibre and sent to the memory experiment (described in the previous section). We have also inserted an electronically controlled mechanical shutter in the signal beam path.

3.1.3 Control

We generate the control field using a second actively mode-locked Ti:Sapph laser (Spectra-Physics Tsunami), pumped by a 5 W 532 nm CW pump. The control laser produces an 80 MHz train of 0.88 GHz bandwidth pulses (~ 500 ps duration) with average power 600 mW. The laser's repetition rate was locked to the signal Ti:Sapph using a commercial lock-to-clock system. The laser's center frequency was locked 6 GHz to the blue of the $6P_{3/2} \rightarrow 6D_{5/2}$ transition using a commercial wavelength meter (see [Appendix B. Laser locking](#) for more information), in order to satisfy the $6S_{1/2} \rightarrow 6D_{5/2}$ two-photon resonance condition with the signal. The spectrum of the pulses directly from the laser, measured with a commercial scanning Fabry-Perót interferometer, is shown in Fig. 3.4.




Figure 3.4: Control Ti:Sapph spectrum, measured on a scanning Fabry-Pérot interferometer (SFP). The visibility of the peaks is limited by mode-matching to the SFP.

The control laser setup is shown in Fig. 3.1 (c). After passing through an optical isolator, the control pulse train is sent into an unbalanced Mach-Zehnder interferometer. There it is split into two, a read-in pulse train and a read-out, both with electronically controlled shutters in their respective beam paths. The read-out is sent, via a polarising beamsplitter (PBS), to a retro-reflector on an electronically controlled translation stage with a ~ 60 cm -long translation range. After passing through a quarter-wave plate (QWP), the pulses are reflected off a mirror and return in the same spatial mode, rotating their polarisation from horizontal to vertical, thus reflecting from the PBS upon return. The folded configuration gives us a total ~ 4 ns range of delay.

The read-in and read-out pulse trains are then overlapped on another PBS and then rotated with a half-wave plate (HWP), passing through a final PBS, in order to project them to a single polarisation, before they are coupled into single-mode fibre and sent to the memory experiment. We are able to independently control the read-in/-out average powers using the HWPs interlaced within the system.

We change the temporal overlap of the read-in control pulse with the signal in the Cs cell by adjusting the delay between the signal and control pulse trains. We accomplish this using the control Ti:Sapph's L2C electronics, which has a built-in fine control of the delay up to 2 ns. For coarse adjustment of delays > 2 ns, we use a discrete variable electronic delay board, which delayed the trigger for the control L2C generated from the signal Ti:Sapph.

3.1.4 Detection

As the signal enters the detection stage of the experiment, shown in Fig. 3.1 (d), it first passes through a 10 nm bandpass filter at 852 nm, in order to filter out any remaining control field leakage due to back-reflections. The signal is then directed towards a Hanbury-Brown Twiss (HBT) setup [165]. The HBT setup is composed out of a HWP and PBS, with two fibre-coupled single-photon avalanche photodiodes (APD) at the outputs (detection efficiency $\eta \approx 50\%$ and dark counts = 296 ± 2 Hz, and $\eta \approx 50\%$ and dark counts = 356 ± 2 Hz, respectively)⁸. The two APDs (with slightly different temporal

⁸Since we use a PBS as the HBT beamsplitter, we could see a slight difference in the measured memory performance depending on which PBS output (i.e. which signal polarisation) we were investigating. We

jitters) are connected to a commercial time-to-digital converter (TDC), triggered from the Pockels cell. For the measurements in this chapter, we sum the clicks from both APDs to determine the total magnitude of the detected signal. For each measurement, we acquired data for 10 s in each of three shutter settings, shown in Fig. 3.5: SIG - both (read-in/-out) control shutters closed, signal shutter open; MEM - all shutters open; and CTRL - signal shutter closed, both control shutters open⁹.




Figure 3.5: Schematic of the different measurement settings. Black rectangles signify a closed shutter.

3.1.5 Evidence of ORCA operation

Fig. 3.6 shows a representative time trace from the time-to-digital converter for the three different measurement settings described in the previous section, proving successful ORCA memory operation for a coherent state input of average photon number $\langle N_{\text{in}} \rangle \approx 2$. Visible on the figure is the read-in bin (at time zero), which contains the input (SIG) and transmitted (MEM) signal, as well as the first three read-out bins¹⁰, at 3.5 ns, 12.5 ns and 16 ns. These arise due to subsequent pulses from the 80 MHz control pulse train reading out residual coherence from the memory. The temporal resolution of the measurement is a convolution of the detector and TDC responses. We define the total memory efficiency as the ratio between the integrated counts in the read-out bin in the MEM setting¹¹ and the integrated counts in the read-in bin in the SIG setting. The storage efficiency on the other hand, i.e. the amount of light absorbed, is equal to one minus the ratio of the MEM and SIG counts in the read-in bin. In the presented example (which was performed for a read-in/-out control pulse energy of $E_{\text{in/out}} = 280$ pJ) the storage efficiency is $\eta_{\text{in}} = 68\%$ and the total efficiency (at 3.5 ns storage time) is $\eta = 17\%$. Immediately visible is the

attribute this to the non-trivial polarisation selection rules when storing in the $6D_{5/2}(F = 2, 3, 4, 5, 6)$ manifold (see [Appendix D. Polarisation](#)).

⁹An additional setting, RI, where the signal and read-in control shutters are open, while the read-out one is closed, is used in chapter 4.

¹⁰The chosen bin size, e.g. for memory efficiency calculation, was 3.2 ns.

¹¹In principle, we should subtract the CTRL counts from the MEM setting, but as we will see in section 3.4, these could also be safely neglected due to the absence of noise in ORCA.

very low noise floor of the memory (CTRL trace in Fig. 3.6).




Figure 3.6: Typical time trace for the ORCA memory for the three different measurement settings: SIG - both (read-in/-out) control shutters closed, signal shutter open; MEM - all shutters open; CTRL - signal shutter closed, control shutters open. Read-in bin (blue shaded area) and first three read-out bins (orange shaded areas) are shown.

In the following sections, we characterise the Cs ORCA memory in more detail, measuring its lifetime, and investigating the scaling of memory efficiency and noise with control pulse energy.

3.2 Memory lifetime

In order to determine the lifetime of the ORCA memory in warm Cs vapour, we measure the memory efficiency for a single-photon level signal (average photon number $\langle N_{\text{in}} \rangle \approx 2$) at different storage times τ , i.e. for different delays between the read-in and read-out control pulses (each with energy of 200 pJ). We determine the storage time by finding

the position of the peak of the retrieved signal pulse, relative to the input pulse, in the measured time traces (such as the one in Fig. 3.6). For storage times $\tau < 5$ ns, we change τ simply by changing the delay in the read-out control arm (see section 3.1.3). We investigate storage times $6 \text{ ns} < \tau < 12.5 \text{ ns}$ in similar fashion, but switching the roles of read-in and read-out using the control Ti:Sapph's L2C electronics. Based on our theoretical estimation of the ORCA memory lifetime in Cs outlined in the next section, this range of storage times was deemed sufficient to measure the lifetime.

3.2.1 ORCA lifetime theory

The two main limiting factors to the ORCA memory lifetime are anticipated to be spontaneous emission from the doubly-excited storage state and motion-induced dephasing due to the wavevector mismatch of the signal and control fields. In the demonstrated case of storage into the $6D_{5/2}$ manifold in Cs, the expected lifetime of the memory due to spontaneous emission would be ~ 51 nanoseconds [166].

Motion-induced dephasing arises due to the Doppler effect and atomic motion in a warm ensemble [145, 167]. This is because the stored excitation is spread over atoms belonging to different velocity classes in the ensemble. Due to the Doppler effect, each velocity class experiences the signal and control fields shifted in frequency. Because the wavelengths of the counter-propagating fields differ significantly in Cs (852 nm and 917 nm), there is only very weak Doppler cancellation. As a consequence, the phase of the stored coherence in different velocity classes will evolve at different rates. Therefore the collective coherence will dephase at a rate proportional to the range of atomic velocities. The motion-induced dephasing lifetime in a thermal ensemble can be estimated as:

$$\tau_D = \frac{1}{k_r v_s}, \quad (3.1)$$

with the shape of the coherence decay being Gaussian [145]. Here, v_s is the average speed of the atoms¹² and $k_r = \frac{2\pi}{\lambda_s} - \frac{2\pi}{\lambda_c}$ is the wavevector of the stored coherence (λ_s and λ_c are the signal and control wavelengths, respectively). In our case, $\lambda_s = 852.35$

¹² $v_s = \sqrt{k_B T / m_{\text{Cs}}}$, k_B being the Boltzmann constant, T the temperature of the atomic vapour and m_{Cs} the atomic mass of caesium.

nm, $\lambda_c = 917.47$ nm, yielding $\tau_D \approx 12.7$ nanoseconds. We therefore expect the ORCA memory lifetime in Cs vapour to be predominantly limited by motion-induced dephasing.

3.2.2 ORCA lifetime measurement

The measured memory efficiency η versus storage time τ is shown in Fig. 3.7. Indeed, the data points seem to follow a Gaussian-shaped inhomogeneous decay curve. We therefore fit a Gaussian function $\eta = \eta_0 e^{-\tau^2/t_{1/e}^2}$ to the data and obtain a $(1/e)$ lifetime $t_{1/e}$ of 5.4 ± 0.1 ns, within an order of magnitude of our estimate of 12.7 ns.




Figure 3.7: The measured total memory efficiency η (orange diamonds) versus storage time τ along with a fit of the memory lifetime curve (blue line) yielding a lifetime $t_{1/e} = 5.4 \pm 0.1$ ns (experimental errors are smaller than the markers).

We attribute the discrepancy to what we call “hyperfine” dephasing, arising from using the whole $6D_{5/2}$ manifold as our storage state. In this case, the stored collective coherence actually consists of a superposition of coherences in the different hyperfine levels of the $6D_{5/2}$ manifold, which are spread over ~ 100 MHz, as shown in Fig. 3.2. These coherences

will evolve at different rates, leading to oscillations in the memory efficiency on a timescale of few ns [18]. Without the presence of motion-induced dephasing one could expect to observe these oscillations, as the different hyperfine coherences rephase. However, due to the significant motion-induced dephasing in warm Cs vapour, the oscillation in our case only appears as additional dephasing. The precise shape of the oscillation will depend on the coefficients in the stored superposition, which in turn depend on the polarisations of the signal and control fields. Hyperfine dephasing can thus be eliminated by appropriate choice of signal and control polarisations, and preparing (via optical pumping) the ground state in the correct Zeeman sub-level [18], so that only one excitation pathway to a single $6D_{5/2}$ hyperfine state is possible.

A memory lifetime of ~ 5 ns gives us a time-bandwidth product of $B \approx 4$. While insufficient for temporal multiplexing applications (see section 1.2.4), it may be enough for certain quantum state manipulation (section 1.2.5). Furthermore, the current memory lifetime could be long enough to make it possible to apply elaborate techniques of extending it, discussed in chapter 6. We also note, that by using a different atomic species with smaller signal/control wavelength mismatch as the storage medium, the memory lifetime can be significantly extended. For example, applying Eq. (3.1) to the $5S_{1/2} \rightarrow 5P_{3/2} \rightarrow 5D_{5/2}$ cascade in rubidium (signal at 780 nm, control at 776 nm), we predict a motion-induced dephasing lifetime of $\tau_D \approx 129$ ns. This would lead to time-bandwidth products > 100 , which start to be useful for temporal multiplexing (section 1.2.4).

3.3 Memory efficiency

In order to verify my theoretical modelling in chapter 2, and better understand the ORCA memory behaviour, we measure how the memory efficiency scales with the control pulse energy for a storage time of 3.5 ns, using weak coherent states ($\langle N_{\text{in}} \rangle \approx 2$) as input. The chosen storage time gives us sufficient temporal separation between the read-in and read-out bins, while still staying safely within the limited memory lifetime $t_{1/e} \approx 5$ ns.

Fig. 3.8 shows the measured efficiencies as a function of the read-in and read-out pulse energies, which were kept equal to each other. On the same figure, we also show curves

generated by the numerical model, based on the theory described in section 2.3. The model was fed with the measured experimental parameters, and the simulated efficiencies were multiplied by an explicit dephasing term of the form determined in section 3.2.2. The only free parameters of the model were the temporal overlap of the signal and control pulses in the Cs cell (which we did not measure explicitly), and the dipole matrix elements for the signal and control transitions. Fitting the latter counteracted the uncertainty in determining the actual control pulse Rabi frequency and the dependence of the coupling strength on the optical field polarisations (as mentioned in section 3.2.2). The fitted parameters are shown in Table 3.1.




Figure 3.8: Measured memory efficiencies (points): storage η_{in} , retrieval η_{out} and total η , as a function of control pulse energy (read-in and read-out). Errorbars are smaller than the markers. Also shown is the result of theoretically modelling the data (solid lines).

We see very good correspondence between theory and experiment. Similarly to our theoretical analysis (section 2.3.2), the maximum storage and retrieval efficiencies occur for different control pulse energies. We see the predicted saturation (at around $\sim 75\%$)

	Fitted	Expected
r1	$1.15a_0$	$4.48a_0$
r2	$2.11a_0$	$2.92a_0$
Δt	-0.44 ns	0

Table 3.1: Fitted parameters from Fig. 3.8, and their expected values. r1 and r2 are the dipole matrix elements for the signal and control transitions, respectively (a_0 is the Bohr radius). Δt is the delay between signal and control (at $\Delta t = 0$ the pulses would meet in the center of the cell). The expected values for r1 and r2 were generated using the FAST (Fast Atomic Spectroscopy from Theory) code for calculating atomic properties from referenced level lifetimes, available from [168].

of the storage efficiency at high control field energies, as the AC Stark shift starts counteracting the increase in coupling strength. Moreover, the retrieval efficiency exhibits a drop as we overdrive the memory interaction, as expected. We reach a maximum total memory efficiency of $\eta = 17 \pm 1\%$. The end-to-end device efficiency, which encompasses the losses between the first and final filters (included), is $\eta_{e2e} \approx 7\%$ - relatively high for an unoptimised¹³ proof-of-principle implementation (see section 1.3.1). We can achieve such high end-to-end efficiency, because ORCA does not require lossy filtering of the signal with etalons, as was the case for e.g. the Raman memory [16].

As mentioned in section 2.3.2, the internal memory efficiency could be improved to $> 50\%$ by heating up the Cs cell, and appropriate choice of control pulse energies and signal/control temporal overlap. Furthermore, end-to-end device efficiencies approaching the internal ORCA memory efficiency should be achievable using low-loss components [18]. One can then expect the ORCA memory to be capable of reaching the end-to-end efficiency requirements for most of the quantum memory applications described in section 1.2. For the time being, however we have verified that our theoretical model outlined in the previous chapter accurately describes the experiment, at least in terms of the scaling of memory efficiency with control pulse energy.

3.4 Memory noise level

In my theoretical considerations I have not explicitly included noise, i.e. the memory producing spurious photons contaminating the signal (see section 1.1.5), since in principle

¹³Improvements in end-to-end efficiency could be made by using less lossy filters, fewer, higher quality mirrors, and anti-reflection coating the vapour cell.

the ORCA memory should be inherently noise-free. First of all, as mentioned in section 2.1.4, the four-wave mixing noise affecting Λ -type memories [139] is entirely absent in a ladder-based memory. Similarly, a ladder-based memory does not suffer from imperfect optical pumping leading to spontaneous Raman scattering [86]. The absence of a buffer gas on the other hand eliminates collision-induced fluorescence [169]. Finally, the counter-propagating signal and control fields differ in wavelength by many nanometres. This allows us to completely remove leakage or scattering of the strong control field into the signal mode using off-the-shelf interference filters, as described in section 3.1.1.

In order to verify whether the ORCA memory is indeed noise-free, we measure the signal detected when the control field is present (the memory is “on”), but with no input signal field, i.e. in the “CTRL” setting shown in Fig. 3.5. We repeat this measurement for a range of control pulse energies (similarly to the previous section), in order to see how any detected noise would scale with the strength of the memory interaction. The results of this measurement are shown in Fig. 3.9¹⁴. As reference, we also plot the average number of input photons and the number of photons retrieved from the memory, based on the measurement from the previous section. We also plot the number of “photons” corresponding to detector dark counts¹⁵.

As we see in Fig. 3.9, the detected noise is not only largely independent of the control pulse energy, it is contained within the measured dark counts of the detectors. Thus it can be fully attributed to the latter. This suggests that indeed the ORCA memory generates no noise itself and attests to the sufficient filtering of the control field in our experiment. The lowest μ_1 , i.e. the noise-to-signal ratio for an input photon number of one (section 1.1.5), we extract from our measurement is $\mu_1 = (3.8 \pm 0.9) \times 10^{-5}$ for a memory efficiency $\eta = 17\%$. Note this is only an upper estimate for the actual μ_1 , since the measured “noise” is upper-bounded by the technical noise (dark counts) of our detectors. To our knowledge, we measure the lowest μ_1 ever reported from an atomic quantum memory, and at least a three orders of magnitude improvement over a cavity-enhanced Raman memory [140].

¹⁴The loss after the memory and detector efficiency have been backed out.

¹⁵In order to measure the detector dark counts, we integrate the counts in a time bin that does not contain the signal or control pulses, but is of similar size. We found this number to be consistent with the detector specifications.




Figure 3.9: Noise measurement (see text for more details). Errorbars are smaller than the markers.

In order to attempt single-photon storage, one requires the heralding efficiency of the photon source, i.e. the probability of having a photon at the input of the memory given a heralding event (see section 4.1.2), to be larger than the μ_1 [19]. Our very low μ_1 therefore encourages us to attempt interfacing the ORCA memory with a standard heralded single-photon source. Such an experiment is the topic of the next chapter.

4

Interfacing the ORCA memory with a single-photon source

In this chapter I describe the main result of my DPhil, which is interfacing a broadband heralded single-photon source with our ORCA quantum memory. In order for an optical memory to be useful for synchronising quantum networks (see section 1.2.3), aside from requirements on efficiency and time-bandwidth product (see section 1.1.1 and the previous chapter), the memory has to be capable of preserving both the quantum statistics and correlations of the stored light. Demonstrating this has been challenging in any memory system, let alone at room temperature and with photons from the type of broadband sources commonly used in quantum optics experiments [16, 11].

This chapter describes the first time broadband single photons have been successfully stored and retrieved from a warm atomic vapour, preserving their quantum statistics. I

will start by introducing the broadband single-photon source we used, which was originally built by (then DPhil student, now) Dr. Patrick Michelberger, and then rebuilt and characterised by Dr. Benjamin Brecht (who also wrote the data acquisition software for the experiment in Python). Afterwards, I will present the results of our attempt (with Dr. Patrick Ledingham) to interface the source with the ORCA memory, as described in the previous chapter, and store and retrieve single GHz-bandwidth photons using ORCA in warm caesium vapour.

4.1 Single-photon source

To generate the single photons for our experiment, we use a 1 GHz bandwidth heralded photon source based on type-II parametric downconversion (PDC) in a periodically-poled potassium titanyl phosphate (ppKTP) waveguide, similar to the one described in [170], which was originally interfaced with a Raman memory in Cs [16]. Here, I describe the theory behind parametric downconversion and define the metrics with which we measure the quantum statistics of the light. This is followed by a description of the experimental details of the waveguide ppKTP source and heralding setup.

4.1.1 Parametric downconversion theory

In parametric downconversion, a high-energy “pump” photon is converted with some probability into a pair of photons with lower energies, “signal” and “idler”, whose spectro-temporal properties are determined by conservation of energy and momentum [171]. In type-II PDC, where the photons are produced in orthogonal polarisations, and thus are distinguishable¹, the quantum state produced is actually a two-mode squeezed vacuum state (in the non-depleted pump limit) [56]:

$$|\Psi\rangle_{PDC} = \exp(\alpha \hat{A}_s^\dagger \hat{A}_i^\dagger + \text{h.c.}) |0\rangle, \quad (4.1)$$

¹This description would also be valid for photons in the same polarisation, but distinguishable, e.g. spectrally.

where $\hat{A}_s^\dagger$ and $\hat{A}_i^\dagger$ are the creation operators for the signal and idler fields², respectively, and α is the squeezing parameter. I can Taylor-expand equation 4.1, obtaining:

$$|\Psi\rangle_{PDC} \approx |0\rangle + \sum_j (\alpha \hat{A}_s^\dagger \hat{A}_i^\dagger + \text{h.c.})^j |0\rangle. \quad (4.2)$$

Looking at Eq. 4.2, it does not seem like the PDC output state resembles a single-photon Fock state $|1\rangle$. Yet, if $\alpha \ll 1$, corresponding e.g. to a low pump power, then I may disregard terms that are higher order in α and which will correspond to multiple pump photons getting converted into PDC pairs at the same time. This highlights the competition between the achievable photon brightness (or photon production rate) and single-photon purity (i.e. how close the produced state is to the ideal $|11\rangle$ state) [172].

In order to eliminate the significant vacuum component of the PDC state, we use heralding. This means detecting one of the photons of the generated pair, let us say the idler (also known as the “herald”), and thus projecting the PDC state onto the $|11\rangle$ component:

$$|\Psi\rangle_{PDC} \approx (\alpha \hat{A}_s^\dagger \hat{B}_i^\dagger + \text{h.c.}) |0\rangle. \quad (4.3)$$

By heralding a photon in the idler mode, we are thus left with a single photon in the signal mode. Including the spectral-temporal properties of the state $|\Psi\rangle_{PDC}$ [151], one gets:

$$|\Psi\rangle_{PDC} \approx \alpha \int d\omega_s d\omega_i F(\omega_s, \omega_i) \hat{a}_s^\dagger(\omega_s) \hat{a}_i^\dagger(\omega_i) |0\rangle, \quad (4.4)$$

where $\hat{a}_s^\dagger(\omega_s)$ and $\hat{a}_i^\dagger(\omega_i)$ are the narrowband creation operators, and $F(\omega_s, \omega_i)$ is the joint spectral amplitude (JSA) of the PDC state. The JSA is defined by conservation of energy and momentum, i.e.

$$F(\omega_s, \omega_i) = \alpha(\omega_s + \omega_i) \phi(\omega_s, \omega_i), \quad (4.5)$$

where $\alpha(\omega_s + \omega_i)$ is the pump envelope function and $\phi(\omega_s, \omega_i) = \text{sinc}[\Delta k(\omega_s, \omega_i)L/2]e^{i\Delta k(\omega_s, \omega_i)L/2}$

²For the time being I ignore the spectral properties of the monochromatic $\hat{A}_s^\dagger$ and $\hat{A}_i^\dagger$ operators.

is the phasematching function. Δk is the momentum mismatch of all three involved fields (signal, idler and pump), i.e. $\Delta k = k_s + k_i - k_p$. In principle, it is non-trivial to simultaneously satisfy energy and momentum conservation in a material for the desired wavelengths. That is why PDC sources often rely on quasi-phasematching for tailoring the photon spectra. Quasi-phasematching can be achieved by periodically poling the material, i.e. modulating its index of refraction in a periodic fashion³. If the poling period is Λ , then $\Delta k \rightarrow \Delta k - \frac{2\pi}{\Lambda}$, allowing one to tune the phasematching function and thus the signal/idler wavelengths. This, along with temperature tuning of the PDC crystal and tuning the pump frequency, is crucial for matching a photon source with an atomic memory.

In principle, the JSA can be very broad ($\sim$ THz) and non-factorable, leading to spectral correlations between the generated photons. In this case, the signal will be projected onto a mixed state upon detection of an idler. The only way to herald pure (and thus useful) single photons from a PDC is to work with a factorable JSA [173]. In our case, we achieve this with spectral filtering⁴, which also has the benefit that it matches the photon bandwidth to the memory bandwidth. While filtering the herald (idler) will project the heralded signal to a specific spectrum, there may still be a significant number of unheralded broadband photons in the signal mode, which may disrupt the measurement (especially of the input field statistics, as the retrieved signal will be filtered by the memory). Therefore, in order to obtain the accurate statistics of photons (see next section) used in the memory experiment, it is beneficial to filter both signal and idler modes.

4.1.2 Photon statistics

The probability of detecting a signal photon from a PDC pair, given a heralding detector click (as shown in Fig. 4.1a), is called the Klyshko efficiency [175], and can be calculated according to:

$$\eta_k = \frac{R_{s,i}}{R_i}, \quad (4.6)$$

³This can be done either during growth of the crystal material or subsequently, e.g. via pulsed electric fields, which locally change the orientation of the crystal.

⁴One can also use a cavity-based source, where photon emission is concentrated to a narrow band at the desired frequency [174].

where $R_{s,i}$ is the detected number of D_i - D_s coincidences and R_i is the number of D_i clicks.

It is often useful to know the probability of the presence of the signal photon at different points of the experiment, given a herald click. We call this probability the heralding efficiency η_h , and it is simply the Klyshko efficiency with backed out losses between the point in question and the detectors (including the efficiency of the detectors themselves, η_{det}), i.e. $\eta_h = \eta_k / \eta_{\text{det}} / \eta_{s,\text{after}}$.




Figure 4.1: **a**, Depiction of the $g^{(1,1)}$ cross-correlation measurement. PDC - source generating pairs of orthogonally polarised photons: signal and idler; PBS - polarising beam-splitter; EXP - experiment; $D_{s/i}$ - signal/idler detector. We can measure the coincidences between the detectors. **b**, Depiction of the $g_h^{(2)}$ autocorrelation measurement using an HBT setup. BS - non-polarising beamsplitter; $D_{s1/s2}$ - signal detectors. Coincidences between all three detectors are measurable.

As mentioned in the previous section, PDC generates two-mode squeezed states, which obey strict photon-number correlations between the two modes. In fact, these correlations are stronger than any correlation that could possibly be achieved with classical, independent sources. The strength of these correlations can be quantified by measuring the normalised Glauber correlation functions [20]. For example, the $g^{(1,1)}$ cross-correlation function between the signal and idler (herald) modes, using the setup shown in Fig. 4.1a. $g^{(1,1)}$ is defined as $p_{si}/p_s p_i$, where p_{si} is the probability for a signal-idler detector coincidence click, and p_s and p_i are the signal and idler detector click probabilities, respectively. Values of $g^{(1,1)} > 2$ signify non-classical correlations⁵.

To calculate $g^{(1,1)}$ from the measurements (section 4.2.1), we use

⁵This is discussed in greater detail in section 4.2.2.

$$g^{(1,1)} = \frac{R_{s,i}}{R_s R_i} R_T, \quad (4.7)$$

where, R_T is the total number of detector trigger events during the whole measurement time and R_s is the number of D_s detector clicks.

In order to quantify the quantum statistics of the heralded signal photons one can measure the heralded auto-correlation defined as $g_h^{(2)} = p_{(s1,s2|i)} / p_{(s1|i)} p_{(s2|i)}$. Here, $p_{(s1,s2|i)}$ is the conditional probability of detecting a coincidence between D_{s1} and D_{s2} given a click in D_i , and $p_{(s1|i)}$ ($p_{(s2|i)}$) is the probability to detect a click in D_{s1} (D_{s2}) given a click in D_i , as defined in Fig. 4.1b. Any $g_h^{(2)} < 1$ verifies non-classical photon-number statistics, with a pure single photon having $g_h^{(2)} = 0$ (see Appendix C. $g^{(2)}$). Because of the higher photon-number components, PDC states will, depending on the pump power, have a $g_h^{(2)} > 0$. We evaluate $g_h^{(2)}$ from the measurements using

$$g_h^{(2)} = \frac{R_{\text{trip}}}{R_{s1,i} R_{s2,i}} R_i, \quad (4.8)$$

where R_{trip} is the number of triple coincidences between D_i , D_{s1} , and D_{s2} , and $R_{s1(2),i}$ is the number of D_i - D_{s1} (D_i - D_{s2}) coincidences.

The benefit of using normalised Glauber correlation functions to quantify the quantum statistical properties of light is that both $g_h^{(2)}$ and $g^{(1,1)}$ are in principle insensitive to experimental losses, which are often difficult to estimate properly and may in fact change over time. In practice, this is not entirely true, since setup losses will determine the data acquisition rates, which will influence our ability to gather enough statistics for a reliable measurement⁶. Moreover, if the losses are too high, the contribution from the technical noise of the detectors (“dark counts”) can influence the measured quantum statistics of light [176].

⁶Low memory efficiency combined with a low experiment repetition rate makes it difficult for example to measure the $g_h^{(2)}$ in spin-wave AFC single-photon experiments [114].

4.1.3 Nonlinear crystal




Figure 4.2: Measured (using a spectrometer attached to a single-photon sensitive camera) PDC spectra (also known as phasematching curves) of the waveguides in group B (b1 to b6) of our ppKTP source, at different crystal temperatures (50 to 67°C). Figures courtesy of Dr. Benjamin Brecht.

We realise our PDC source in a periodically-poled, rubidium-exchanged waveguide in a potassium titanyl phosphate (KTP) crystal, commercially available from AdvR. Waveguide systems offer a well defined, collinear spatial mode [177], and high brightness at relatively low pump powers, owing to the strong confinement of light over relatively long

distances and the discretisation of the transverse spatial degree of freedom⁷. We chose KTP as the nonlinear medium, as it has a large nonlinear coefficient and high damage threshold at <NIR wavelengths [179].

The chip we used was 2 cm long and hosted three groups of waveguides (A, B, C), each containing 6 guides, spaced by $35 \mu\text{m}$. The nominal widths were: 2, 3, 4, 2, 3, $4 \mu\text{m}$, designed to sustain single-mode operation at 852 nm ⁸. The crystal was periodically poled with a poling period $\Lambda \approx 10 \mu\text{m}$. By appropriately choosing the guide (different guides have slightly different phasematching and thus generate PDC a different wavelengths, see Fig. 4.2) and adjusting the temperature of the crystal, we could tune the phasematching of the source to obtain roughly degenerate photon pairs at 852 nm , as shown in Fig. 4.3.




Figure 4.3: Extracted (from Fig. 4.2) wavelengths of signal and idler emission peaks for different PDC crystal temperatures. Shown also are linear fits to the data, and the point of degeneracy, i.e. where signal and idler photons have the same wavelength. Figure courtesy of Dr. Benjamin Brecht.

⁷In fact, waveguide sources are typically at least an order of magnitude brighter than bulk crystal sources [178].

⁸This condition does not hold true for the $4 \mu\text{m}$ wide waveguides, as seen from the multi-peak structures in the lower panel of Fig. 4.2.

4.1.4 Crystal mount

The PDC waveguide chip source was mounted on a homebuilt copper⁹ mount (Fig. 4.4). The mount rested on a Peltier heating element, which was glued with thermal paste to a thick aluminum lever¹⁰ which acted as a thermal load. We also inserted a heat cartridge into the lever in order to decrease the temperature difference between the latter and the copper mount, improving the operation of the Peltier element. The temperature of the copper mount was controlled to 0.1°C degree accuracy using a commercial temperature controller (Newport 350B) connected to the Peltier element, and monitored with a temperature sensor inserted into the mount.




Figure 4.4: ppKTP mount (copper, orange) and thermal load (aluminum, grey). The ends of the hollow channel along the crystal groove were capped in the final device.

The ppKTP crystal was inserted into a groove in the copper mount, on the bottom of which we made narrow holes connected to a hollow channel inside the mount. Said channel was capped at both ends and, in the middle, was connected to a tap on a third side of the mount. This tap was connected to a commercial fishtank pump, in which we reversed the direction of air flow. This generated a weak vacuum that sucked the ppKTP

⁹Copper was chosen for its good thermal conductance.

¹⁰The lever was designed so as to be easily exchangeable with the crystal mount used in [16].

crystal in place, stabilising its position.

4.1.5 Experimental setup




Figure 4.5: Experimental scheme for the heralded single-photon source. SHG - bulk periodically-poled potassium titanyl phosphate (ppKTP) crystal; DM - dichroic mirror; ppKTPw - ppKTP waveguide; EF - edge filters; HWP - half-wave plate; PBS - polarising beamsplitter; FP - Fabry-Pérot etalons; APD - single-photon avalanche photodiode; Exp - single-mode fibre leading to the memory experiment (see section 3.1.1).

The experimental setup for generating the heralded single photons is shown in Fig. 4.5. We used one of the $3\ \mu\text{m}$ wide waveguides from the B group to generate our photon pairs (see section 4.1.3), chosen due to it being the waveguide with the highest brightness and thus photon-pair production rate, operating it around the degeneracy point, i.e. around 70°C ¹¹. The waveguide is pumped by 426 nm pulses in a single-pass configuration. The pump light is produced by frequency doubling the 80 MHz pulse train from the 852 nm Ti:Sapph (see section 3.1.2) via second harmonic generation in a 2 mm long ppKTP bulk crystal. At an incident average power nearing 700 mW we generate $\sim 4\text{ mW}$ of average power at 426 nm before the PDC waveguide in-coupling lens. The fundamental light is separated from the second-harmonic using a short-pass dichroic mirror, and sent into the weak coherent state signal setup (Fig. 3.1 (b) in section 3.1). The second-harmonic pump is coupled to the PDC waveguide with a total transmission of $< 10\%$ including the losses at the in- and out-coupling lenses. The waveguide was not designed to be single-mode for the pump wavelength, and we optimise the coupling of the pump to primarily excite the zeroth-order spatial mode, hence the low overall transmission. The frequency-degenerate,

¹¹The precise operating temperature would depend on the current environmental conditions.

polarisation-orthogonal PDC signal and idler photons have a bandwidth on the order of THz. After the ppKTP waveguide, the signal and idler are separated from the pump using another short-pass dichroic mirror.




Figure 4.6: Blue - measured transmission spectrum of the filtering setup from Fig. 4.5. Orange - fitted Lorentzian curve, giving a FWHM of 0.95 GHz. The center of the transmission peak is located 6 GHz to the red of the $6S_{1/2}(F = 4) \rightarrow 6P_{3/2}$ transition, i.e. at the center frequency of the weak coherent state signal from section 3.1.2.

After passing through the dichroic mirror, we filter the THz-bandwidth signal and idler fields, in order to match them to the atomic memory. First, we apply coarse filtering using a set of two edge filters, a short-pass and a long-pass, to get rid of any higher order mode leakage¹². Then, the modes are spatially separated via a polarising beamsplitter (PBS) and coupled to their own single-mode fibres (SMF) with an efficiency of $(64 \pm 1)\%$ for the signal mode and $(53 \pm 2)\%$ for the idler mode¹³. The modes are then out-coupled and recombined on a PBS, forming a common spatial mode that passes two etalons, one of 18.4 GHz free spectral range (FSR) and one of 103 GHz (giving an effective FSR of the combined system ~ 1 THz). This is followed by a holographic volume Bragg grating

¹²Higher order spatial modes of the signal/idler photons in the waveguide will have different phase-matching, and thus the photons generated will be shifted in wavelength by several nm. These higher order modes could be seen in the phasematching spectra in Fig. 4.2, section 4.1.3, as the small side peaks.

¹³We use a classical 852 nm “alignment” mode (which is coupled to the zeroth order mode of the waveguide and thereby comparable to the signal and idler modes) to characterise beam propagation transmission. This allows us to characterise the modes with standard power meters and photodiodes.

with a width ~ 100 GHz, to get rid off any residual light unfiltered by the etalons. Using a narrowband ($\sim$ MHz) CW probe beam we measure the total filtering scheme's width to be ~ 1 GHz for both modes, as shown in Fig. 4.6. This defines the bandwidth and center frequency of the single photons we intend to store in the memory.

After the filtering, the modes are spatially separated again via a PBS. The idler is then coupled to a single-photon avalanche photodiode (efficiency $\eta \approx 50\%$, dark counts = 163 ± 1 Hz) via a multimode fibre. The total transmission from after the waveguide to in front of the idler detector is $\eta_{i,\text{filt}} = (9.7 \pm 0.1)\%$. The signal, on the other hand, is coupled to an SMF (total transmission from after waveguide to after this SMF is $\eta_{s,\text{filt}} = (12.8 \pm 0.3)\%$), and sent toward the ORCA memory using the weak coherent state signal mode from chapter 3. The total transmission from the PDC source to in front of the signal detectors (averaging over the two detector SMF couplings) is $\eta_{s,\text{total}} = 3.7 \pm 0.1\%$. Thus, the signal photon will experience an additional $\eta_{s,\text{add}} \approx 30\%$ transmission after the filtering stage.

To obtain the results presented in the rest of this chapter, we pumped the PDC source with 4 mW of average power in front of the waveguide in-coupling lens. This gives us an idler (signal) count rate of around 30 kHz (10 kHz) with the memory off. The typical Klyshko efficiency η_k measured is 0.7 %. Including the losses mentioned above, this yields a heralding efficiency of $\eta_h = 4.8\%$ ¹⁴, which is well above the $\mu_1 \approx 10^{-5}$ of the Cs ORCA memory (section 3.4), in principle enabling the storage and retrieval of single photons.

4.2 Interfacing PDC source with ORCA in Cs

We now have a heralded photon source which should be roughly matched in center frequency and bandwidth to the ORCA quantum memory, as described in section 3.1.1. We couple the PDC signal photons to the memory using the weak coherent state signal mode (section 3.1.1), and then detect them after the memory using a Hanbury-Brown Twiss type setup (section 3.1.4), composed of a beamsplitter and two single-photon de-

¹⁴As a cross-check, we calculate the heralding efficiency just after the waveguide to be $\eta_{s,\text{waveguide}} = 38\%$. This should in principle be 100%, and we attribute the missing factor of 2.6 to not measuring explicitly the loss inside the waveguide, the out-coupling loss from waveguide to free-space, and the potential frequency mismatch of the etalon centroids between signal and idler.

tectors connected to a time-to-digital converter, similarly to the idler detector. Assuming the memory is noise-free (section 3.4), we can optimise the memory efficiency by simply maximising the raw counts on the two signal detectors in the read-out time bin (more on this in the next section). In this way, we optimise the efficiency by adjusting the control field polarisation, relative delay¹⁵, and read-in/-out pulse energies. The memory efficiency was maximised for read-in/-out pulse energies of 0.21(1)/0.97(1) nJ, and we operated the memory at a storage time of 3.5 ns (same as for the weak coherent state measurements in sections 3.3 and 3.4). The memory efficiency was quite insensitive to the center frequency of the control field¹⁶, possibly due to the PDC signal having a larger bandwidth (0.95 GHz) than the control (0.88 GHz).

4.2.1 Data collection

Similarly to the weak coherent state measurements (chapter 3), we use three mechanical shutters (see Fig. 3.1) to define the different measurement configurations: “input” signal measurements with signal shutter open, and both read-in and read-out closed (SIG); storage measurements with signal and read-in shutters open, and read-out shutter closed (RI); full memory (storage and retrieval) measurements with all three (signal, read-in and read-out control) shutters open (MEM); and noise measurements with read-in and read-out control shutters open, and signal shutter closed (CTRL). A single measurement run consisted of recording the number of detector counts in the MEM configuration for 180 s, followed by recording the counts over 10 s in the respective RI, SIG, and CTRL configurations. Upon completion of one run, the registered data was written to disk and the measurement repeated. We decided to operate the experiment in this interspersed way, in order to minimise the effects of slow drifts in the setup¹⁷. We acquired data for over 90 hours in one continuous measurement run, sporadically peaking up the coupling of the pump light into the PDC source, and coupling of the signal and idler fibres after the PDC

¹⁵The setup was built such that the PDC photons would be roughly temporally overlapped with the weak coherent state signal. Thus only fine tuning of the L2C delay was required to temporally overlap the former with the read-in control pulses.

¹⁶Interestingly, the memory efficiency for PDC photons was highest, when the control field center frequency was red-shifted by 200 MHz with respect to the assumed two-photon resonance condition. This might have as well been due to the imprecise estimation of the signal filter center frequency however.

¹⁷These could be caused by changes in the laboratory temperature and humidity.

source, in order to mitigate long-term drifts in alignment. However, after around 30 hours the setup still drifted too strongly, thus we limit our analysis only to data acquired in the first ~ 30 hours of the measurement (corresponding to 521 measurement runs).

For each shutter configuration in each measurement run, using the time-to-digital converter (TDC), we recorded arrival time histograms for the three detectors, i.e. the idler APD (as shown in Fig. 4.5), which we call D_i , as shown in Fig. 4.1, and the two signal APDs (shown in Fig. 3.1 (d)), henceforth known as D_{s1} and D_{s2} (Fig. 4.1b). More specifically, these are histograms of APD detection times with respect to a 1 MHz trigger signal derived from the 852 nm Ti:Sapph. The histogram time bin width for these “singles” measurements was 200 ps¹⁸.

In order to compare different measurement configurations with potentially different total measurement times, we add all arrival time histograms for a given measurement configuration and normalise them to both the number of measurements (521) and the respective measurement duration (180 s for MEM, 10 s for the other configurations), obtaining a countrate per time bin in units of Hz. Fig. 4.7 shows the resulting accumulated arrival time histograms for D_i , D_{s1} , and D_{s2} (a to c) in the SIG (blue trace) and MEM (orange trace) shutter configurations. Note that these are raw data (without e.g. any background subtraction or heralding), which demonstrates the high signal-to-noise ratio of the measurements, signified by the good visibility of the peaks (see especially Fig. 4.7d). Furthermore, as one can see in the same figure, due to the short memory lifetime (section 3.2), and without the need for time-consuming atomic state preparation (as the storage state is spontaneously empty even at high temperatures), we are able to operate the experiment at the full 80 MHz repetition rate of our PDC pump.

To reduce the impact of noise counts (predominantly from detector dark counts, as was discussed in section 3.4), we applied time gating to the recorded arrival time histograms and only used detection events that lay within the time gates in our analysis. The read-in time gates (blue regions in Fig. 4.7d) are centred around the read-in peaks and have a width of 2.5 ns¹⁹. The read-out time gates (orange regions in Fig. 4.7), of similar width,

¹⁸This was chosen as a compromise between temporal resolution of the TDC and total number of time bins in the histogram.

¹⁹Chosen so that the peaks were completely inside the gated region and did not overlap with the read-out bin.




Figure 4.7: **a**, Histogram of click times of detector D_i with respect to a 1 MHz trigger. **b**, The same histogram for detector D_{s1} , for SIG (blue trace) and MEM (orange trace) shutter configuration. The reduced height of the main peaks in the MEM configuration signifies read-in, i.e. storage of photons. The small peaks 3.5 ns after the main peaks originate from light that has been retrieved from the memory (more visible in **d**). **c**, The same as in **b**, for detector D_{s2} . **d**, Section of **b** with additional CTRL configuration data (purple). The shaded regions indicate the time gates.




Figure 4.8: **a**, Accumulated normalised correlation histogram for two-fold coincidences between detectors D_i & D_{s1} , shown for the SIG (blue trace) and MEM (orange trace) configurations. Note the logarithmic scaling of the y axis. Shown shaded are time gates used for the data analysis. Successive read-in (read-out) time bins are visualised by blue (orange) shaded regions with decreasing saturation. **b**, The same as in **a**, however for two-fold coincidences between detectors D_i & D_{s2} . **c**, Correlation histogram for three-fold coincidences between detectors D_i & D_{s1} & D_{s2} .

were offset from the corresponding read-in time gates by 3.5 ns (the memory storage time chosen for the experiment). In order to calculate the memory efficiency we integrate the detection events over the gated regions (similarly as in section 3.1.4), .

In addition to “singles” histograms described above, we also simultaneously recorded coincidence histograms for signal-idler two-fold coincidences (D_i - D_{s1} , D_i - D_{s2}) and the three-fold coincidences (D_i - D_{s1} - D_{s2}). These are start-stop histograms, where the detection of an idler photon starts the measurement and the detection of a signal photon (or the detection of a D_{s1} - D_{s2} coincidence) serves as the stop signal for the two-fold (or three-fold) coincidence measurement, yielding time differences between the two detection events. For these measurements, the time bin width of the TDC was chosen to be 100 ps²⁰ and the time gates had a width of 3.5 ns²¹. Again, data was recorded for each shutter configuration

²⁰To ensure that the temporal resolution of the measurement was not limited by the TDC.

²¹Since the pulses would appear longer in the coincidences measurement due to the convolved temporal response of two detectors, instead of just one for the “singles” measurement.

in each measurement run, and then post-processed in a similar way as the “singles” histograms for visualisation (Fig. 4.11), yielding coincidence rates per time bin in units of Hz.

The SIG traces in Fig. 4.11 exhibit a main peak at a time difference of 0 ns, with subsequent smaller peaks at integer multiples of 12.5 ns (inverse of the 80 MHz pump repetition rate). The ratio of the area under the main peak (the number of counts within the read-in time gate) and the area of a side peak (the number of counts within a time gate 12.5 ns away), or the so-called coincidences-to-accidentals ratio (CAR), in our case amounts to ~ 130 . This number is also a measure of the amount of higher-order photon-number components²² in our two-mode squeezed PDC state. The main peak originates from correlated idler and input signal photons, which have been generated in the same PDC event. Therefore, if the probability to generate a photon pair is given by p , we expect the area under the main peak to be proportional to p . In contrast, the other peaks (in the SIG configuration) originate from two uncorrelated PDC events, where signal and idler photons have been generated by subsequent pump pulses. Consequently, we expect that the area of a side peak would be proportional to p^2 . From this we estimate $p^2/p = p \approx 7.7 \cdot 10^{-3} \approx 1/130$, which puts us into the weakly pumped regime of PDC (see section 4.1.1) or, in other words, confirms that we primarily generate pairs of single photons during the PDC²³.

4.2.2 Results

Here I present the results of interfacing the ppKTP heralded single-photon source with the Cs ORCA memory (described in section 3.1.1). We first verify that the PDC source is frequency and bandwidth matched to the memory, by showing that we are able to store PDC photons in the ORCA memory as efficiently as weak coherent states. We then verify quantum ORCA operation by measuring the two normalised Glauber correlation functions, the signal-idler cross-correlation $g^{(1,1)}$ and heralded autocorrelation $g_h^{(2)}$ of the signal, at the memory input and output.

²²Proportional to the squeezing parameter α^2 , as defined in section 4.1.1.

²³Since the probability to generate two photon pairs in one single PDC event is approximately given by $p^2 \approx 6 \times 10^{-5}$, i.e. is negligible.


Memory efficiency

Figure 4.9: Change in total internal memory efficiency η over measurement time.

We estimate the ORCA memory efficiency for PDC photons by calculating the ratio between the total number of signal-idler coincidences in the MEM read-out bin and SIG read-in bin (Fig. 4.8). Fig. 4.9 shows how the memory efficiency changed over time. The average efficiency over the measurement time was $\eta = 15 \pm 2 \%$. This is slightly lower than the maximum efficiency measured for weak coherent states (17%, section 3.3), though we note that the PDC memory efficiency was not systematically optimised (which explains its slow drift over time), and there is a larger bandwidth mismatch between the 0.88 GHz control field and ~ 1 GHz PDC photons (compared to the 0.8 GHz weak coherent signal)²⁴. It is nonetheless impressive that the memory efficiency remained quite stable and above 10% for over 30 hours of continuous measurement, with practically no supervision required²⁵. This level of robustness is necessary if the device is to be used reliably in practical quantum networks.

²⁴This is especially visible in Fig. 4.11, in how the retrieved photons are noticeably longer in time, due to being re-shaped upon read-out by the lower bandwidth control field. A similar effect was observed in the Raman memory in [66].

²⁵Due to the extensive filtering system based on Fabry-Perot etalons, the Raman memory required re-optimisation every 30 minutes or so over the many hour data acquisition times [16].

Signal-idler cross-correlation

To verify the quantum performance of the memory we measure the photon number statistics and correlations of the retrieved signal, and compare it with the input. First, as a measure of the quantum correlations between the signal and idler modes, we evaluate the second-order cross-correlation function $g^{(1,1)}$, using Eq. 4.7 and the data presented in Fig. 4.7 and 4.8. The results for the SIG configuration are summarised in the first row of Tab. 4.1, where we find $g^{(1,1)} = 131(1)$ for a time difference of 0 ns and $g^{(1,1)} \approx 1$ for integer multiples of 12.5 ns (since the photons detected there mostly come from separate PDC events). We also note that the values at the read-out times (3.5 ns offset from the 12.5 ns time slots) are meaningless, since we expect no actual signal at the detectors.

Similarly to the SIG setting, in the MEM configuration in Fig. 4.8, we find a dominating peak in the correlation histograms at a time difference of 0 ns, coming from photons transmitted by the memory. In addition, we see side peaks at integer multiples of 12.5 ns²⁶, as well as peaks that are offset from these major peaks by 3.5 ns. These originate from coincidence events between idler photons and signal photons that have been stored in and retrieved from the memory. The reason for this series of peaks lies in the non-unity retrieval efficiency of our memory (similarly to Fig. 3.5 in section 3.1.4). A stored photon is not necessarily read out after 3.5 ns, but with some probability (given by one minus the retrieval efficiency) may remain stored in the memory. Then, it can be read out by subsequent pulses from the control pulse train.

Similarly to the SIG setting, we evaluate $g^{(1,1)}$ in the read-in and consecutive read-out bins in the MEM setting. The results are summarised in the second row of Tab. 4.1. As one can see, we find that the $g^{(1,1)}$ remains non-classical (> 2) for up to 16 ns, i.e. around three times the lifetime of our memory (section 3.2.2). In other words, non-classical photon correlations are retained even after the memory efficiency has dropped to around 5% of its initial value. We attribute the reduction of the $g^{(1,1)}$ in consecutive read-out bins to increased relative dark count contamination of the weak signal [176].

Focussing on the first read-out bin (at 3.5 ns), we measure $g^{(1,1)} = 120 \pm 5$, equal within experimental error to the input signal value of $g^{(1,1)} = 130 \pm 5$. We thus demonstrate

²⁶Note the side peak at 12.5 ns being higher than the corresponding peak for the SIG configuration.

	$\tau = 0$ ns	$\tau = 3.5$ ns	$\tau = 12.5$ ns	$\tau = 16$ ns	$\tau = 25$ ns	$\tau = 28.5$ ns
SIG	130(5)	—	1.1(2)	—	1.2(2)	—
MEM	80(3)	120(5)	9.7(4)	11.3(5)	1.7(2)	1.1(1)

Table 4.1: Calculated signal-idler cross-correlation $g^{(1,1)}$ for consecutive time-bins (see Fig. 4.8) with the memory off (SIG) and on (MEM).

the preservation of non-classical photon-number correlations in ORCA²⁷. This is better seen in Fig. 4.10, in which we show how the $g^{(1,1)}$ in this bin changed over time. As one can see, within experimental error, the $g^{(1,1)}$ of the retrieved field followed closely the input signal²⁸, again demonstrating the excellent robustness of the ORCA memory. The significant drift over time in the $g^{(1,1)}$ values themselves is mainly caused by the waveguide in-coupled PDC pump power changing over time due to e.g. temperature fluctuations.




Figure 4.10: Change in the value of the signal-idler cross-correlation $g^{(1,1)}$ over measurement time for the input field (Input) and the read-out field (Memory). Also depicted is the classical bound.

Often in quantum memory literature, in order to verify the preservation of non-classical correlations of the stored light, the Cauchy-Schwarz inequality is tested. In other words, an optical memory is shown to be capable of storing quantum correlations,

²⁷Again, we can attribute the slight reduction in the mean value of the $g^{(1,1)}$ in the read-out to dark count contamination [176].

²⁸The drift of the latter was due to the coupling of the pump into the PDC source drifting over time, since the $g^{(1,1)}$ is actually a function of the mean photon number of the PDC state [180], which, in turn, depends on the in-coupled pump power (as discussed in section 4.1.1).

if the Cauchy-Schwarz parameter $R = \frac{(g_{s,i}^{(2)})^2}{g_{s,s}^{(2)} g_{i,i}^{(2)}} > 1$ [42]. Here, $g_{s,s/i,i}^{(2)}$ is the second-order signal-signal/idler-idler unheralded autocorrelation, and $g_{s,i}^{(2)} = g^{(1,1)}$. In our measurement we did not have access to the $g_{s,s/i,i}^{(2)}$ correlations, but we can upper-bound them by $g_{s,s/i,i}^{(2)} \leq 2$ ²⁹. The estimated Cauchy-Schwarz parameter for our retrieved field is therefore $R \geq 3600$, i.e. significantly above the classical bound of 1 and at least two orders of magnitude larger than measured in a recent experiment interfacing a heralded PDC source with an AFC quantum memory [114]. This proves ORCA's ability to preserve non-classical photon-number correlations.

²⁹The highest value of $g_{s,s/i,i}^{(2)}$ possible is 2, which corresponds to a single-mode thermal state, which is to be expected from half of a pure two-mode squeezed state of light. In reality though, $g_{s,s/i,i}^{(2)} < 2$, as the thermal state will usually have some degree of multimodeness stemming from spectral correlations in the original PDC-generated squeezed state (section 4.1.1).

Heralded autocorrelation of the signal




Figure 4.11: **a**, Input signal field: orange - histogram of the time difference Δt between idler detector clicks, and signal detector coincidences, i.e. triple coincidence histogram R_{trip} ; blue - product of the two-fold coincidences between the idler and two signal detectors, $R_{s1,i}$ and $R_{s2,i}$, normalized by the idler counts R_i , i.e. predicted triple coincidence histogram for coherent light of the same average photon rate as the PDC. **b**, Same as **a**, but for the retrieved field.

As an even stronger benchmark of quantum performance, we evaluate the second-order heralded auto-correlation function $g_h^{(2)}$, using Eq. 4.8 and the data shown in Fig. 4.11. It is due to the low rate of triple coincidences (signifying good quality single photons are being produced by our PDC source and being retrieved from the memory), that our data acquisition time had to be on the order of days in order to gather acceptable statistics. Comparing the orange and blue histograms in Fig. 4.11 **a/b**, we clearly see the anti-bunched nature of the input/retrieved field, with the measured triple coincidences (orange) being much lower than what one would expect from a coherent state (blue). The ratio between the orange and blue histograms is in fact the heralded auto-correlation

function $g_h^{(2)}$. By integrating the histograms in panel **a** of the figure over the read-in time bin, we calculate the $g_h^{(2)}$ of our input field to be 0.02 ± 0.01 , well below the classical boundary of 1 and very close to the ideal single-photon value of 0, as expected from weakly pumped PDC (see section 4.1.2).

We then perform the same calculation for the read-out bin in the MEM configuration, and we find the heralded auto-correlation function of the retrieved light to be $g_h^{(2)} = 0.03 \pm 0.01$, i.e. within experimental error of the input value³⁰. In fact, this is the lowest $g_h^{(2)}$ ever measured from a room-temperature quantum memory, and taking into account experimental error, the lowest bounded $g_h^{(2)}$ measured from any memory system to-date. The excellent agreement between input and output $g_h^{(2)}$ further confirms the memory generates no noise (section 3.4). The quantum performance of ORCA is on par with a fundamentally noise-free (non-atomic) memory based on electronically switched free-space cavities [11], though due to the technical simplicity of the former, ORCA should be much easier to integrate, which is the topic of the next chapter.

³⁰The slight increase in the mean value of $g_h^{(2)}$ of the retrieved field can be attributed to a stronger contribution of detector dark counts in the weak signal [176].

5

Cs-filled hollow-core photonic crystal fibres as an integrated platform for quantum memories

In order to be efficient, quantum memories require strong light-matter coupling to facilitate the transformation of a photon into a matter excitation. We have seen in the theoretical analysis in chapter 2, that the ORCA coupling strength can be increased from the material side, by increasing the number of atoms interacting with light (parametrised by the optical depth, OD), or from the photonic side, by increasing the Rabi frequency of the control field. Aside from increasing the control pulse energy, the latter can be accomplished by confining the light in a waveguide geometry, enabling tighter focussing and increased light intensity, while maintaining a long interaction length [15]. Furthermore, a waveguide-based memory is desirable, as it makes the device potentially easier to interface with integrated photonic circuits.

Warm vapours contained in hollow-core photonic crystal fibres (HC-PCF) [181] would be a natural platform for integrating the room-temperature ORCA memory, described in the previous chapters. HC-PCFs filled with alkali vapor have already been used to demonstrate the basic building blocks required for photonic quantum information processing, i.e. single-photon level optical nonlinearities [182] (with the potential of using Rydberg atoms [183]), and even the Raman memory [15]. However, due to the alkali's reactivity, adsorption to the fibre core surfaces significantly limited the ambient optical depths achievable in HC-PCFs, even at high temperatures [184]. To overcome this limitation, one can utilise the phenomenon of light-induced atomic desorption (LIAD), in which a strong off-resonant beam evaporates the adsorbed alkali, increasing the atomic density in the vapour phase [185].

This chapter is devoted to the progress we have made on a caesium-filled HC-PCF platform towards fibre-integrating a room-temperature quantum memory, following the pioneering work of former DPhil student Michael Sprague [184, 15]. I will describe the kagomé-type HC-PCFs (drawn for us by the group of Philip Russell at the Max Planck Institute for the Science of Light in Erlangen, Germany) and ultrahigh vacuum chamber (originally designed and constructed by Michael Sprague, and then, due to a leak, dismantled, cleaned and rebuilt by myself) that are at the heart of the experiment, and then explain how I used near-infrared LIAD to generate record high optical depths of caesium in the kagomé-type HC-PCFs. I will then discuss my efforts at modelling the LIAD phenomenon, and briefly touch on our (with Dr. Amir Feizpour) attempts at using green light for the desorbing beam.

5.1 Cs-filled hollow-core fibres

Here I describe the Cs-filled hollow-core fibre platform designed for implementing a warm vapour quantum memory. I start by introducing the type of hollow-core photonic crystal fibre used in our experiments, the kagomé-structured fibre, after which I outline how we filled our fibres with room-temperature Cs vapour within an ultra-high vacuum system. Our hope was that a fibre with a large core size of $46\text{ }\mu\text{m}$ could support sufficient ambient optical depths of Cs to be able to implement a quantum memory.

5.1.1 Kagomé-structured hollow-core fibres

Traditionally, solid-core glass fibres have been used to guide light via total internal reflection in a core of higher index of refraction than the cladding [186]. In hollow-core fibres, however the guiding mechanism usually involves a form of photonic bandgap that decouples the modes in the fibre core from the cladding modes [187]. This is especially visible in kagomé-structured fibres [188], where due to the very small struts forming the microstructured cladding (as shown in Fig. 5.1), the cladding modes have very low overlap with the core modes [189]. This can be seen in Fig. 5.2, where we selectively couple light into the core and cladding separately. This decoupling of the core and cladding modes prevents the leakage of light from the core to the cladding, effectively containing the light in the core and thus guiding it.

Kagomé-structured fibres possess a number of advantages towards the quantum memory application we have in mind. First of all, they exhibit very broad guidance, as shown in Fig. 5.1 (d). This is essential for at least two reasons. It makes it possible to implement the ORCA protocol, in which usually signal and control fields have quite different wavelengths. Secondly, it allows us to use a large range of LIAD (see section 5.2) wavelengths, which have to be far away from any atomic resonances (so as not to influence the memory process directly). Furthermore, kagomé fibres can support relatively large core sizes ($> 10 \mu\text{m}$) [188], leading to extended transit times of atoms moving through the core. Atomic transit time will ultimately limit the memory lifetime in hollow-core fibres, since atomic collisions with the core walls are in principle neither spin nor orbital preserving, and thus will lead to dephasing of the memory [15]. For spin-based memories one could use spin-preserving coatings [190, 191]. Their use in confined geometries however, is still under investigation [192]. For the orbital-based ORCA memory, the lifetime of which is expected to reach around a hundred nanoseconds in Rb vapour (see section 3.2), fibre core sizes greater than $\sim 20 \mu\text{m}$ should be sufficient so as not to limit the memory lifetime. The large cores of kagomé fibres also facilitate loading the fibres with warm vapour, as we will see in section 5.1.2.




Figure 5.1: (a) SEM image of the 26 μm core kagomé-structured hollow-core photonic crystal fibre. (b) Same as (a) for the 46 μm HC-PCF. (c) Optical image of the fundamental mode of the 46 μm HC-PCF. (d) Loss spectrum of 26 μm HC-PCF. All images courtesy of Amir Abdolvand from the Max Planck Institute for the Science of Light in Erlangen.




Figure 5.2: (a) Optical image of the front face of the 26 μm core kagomé-type HC-PCF with light coupled into the fibre core. (b) Optical image of the front face of the 46 μm core kagomé-type HC-PCF with light coupled into the fibre cladding.

In the experiments described in this chapter, we have used two kagomé-type hollow-core photonic crystal fibres of 26 μm and 46 μm core diameter (Fig. 5.1 (a) and (b),

respectively). Two fibres of different core size were initially chosen in order to compare the Cs vapour filling time and achievable ambient optical depths (see next section), and the transit time of Cs atoms in the core. The losses at 852 nm in the 26 μm and 46 μm fibres were 1.5 dB/m and 2.6 dB/m, respectively. While relatively high for fibres¹, at the short lengths (~ 20 cm) we have used, this loss would not significantly affect our experiments. Our fibres were in principle spatially multimode at 852 nm, though with careful coupling, it was possible to excite predominantly the fundamental mode, as shown in Fig. 5.1 (c). Interestingly, while not designed to be polarisation maintaining, our fibres turned out to preserve polarisation quite well² (using cross-polarizers, we measured polarisation extinction ratios above 400:1).

5.1.2 Filling the HC-PCFs with Cs

Caesium, as a heavy alkali metal, reacts strongly (explosively) with oxygen and water, hence it has to be contained in an air-free environment. In order then to fill our fibres with Cs vapour and perform experiments, we use a home-built ultra-high vacuum system, as shown in Fig. 5.3 [193]. The vacuum system was designed to hold fibres of ~ 20 cm length, with optical access at both ends to couple light in and out of the fibres. The fibre ends were held next to each other in the v-grooves of custom-built fibre holders inside the chamber. In order to connect the different vacuum components together, we used conflat flanges, i.e. knife-edge flanges with soft metal gaskets, for optimal vacuum performance³. The vacuum chamber included an inlet to introduce nitrogen gas⁴, though it was never used.

The vacuum system was pumped by a combination of roughing pump and turbomolecular pump, allowing us to reach a base pressure of around 10^{-10} Torr, as measured on an ion gauge integrated into the system⁵. Such low pressure was achieved by wrapping

¹Modern telecommunications fibre has loss in the infrared on the order of 0.2 dB/km [36].

²We have also noticed that depending on the input light polarisation, one could selectively couple to different higher order spatial modes.

³Unlike seals based on rubber o-rings, copper gaskets can be baked at high temperatures, and would not react with Cs to such an extent, preventing leaks.

⁴Since originally the system was designed for the Raman memory, in which case nitrogen buffer gas improves optical pumping and atomic transit time [100].

⁵The ion gauge failed after 6 months though, probably due to Cs accumulation on its electrodes, and so we could only monitor the quality of vacuum using Pirani gauges with a lower range of 10^{-4} mbar.

the system in heating wire and aluminum foil, and baking it at $\sim 200^\circ\text{C}$ for 2 days. It is important to ramp the temperature up and down gradually, so as not to create stress in the windows that could lead to leaks. A pneumatic gate valve was installed after the turbomolecular pump. In case the electric power to the pumps was cut, the valve would close, maintaining the vacuum in the experiment⁶.




Figure 5.3: Model of the ultra-high vacuum system with legend. Image courtesy of Michael Sprague.

Caesium was introduced into the vacuum system by inserting a caesium ampoule into a flexible bellows in the chamber (marked Cesium reservoir in Fig. 5.3). The ampoule was consequently broken by bending the bellows. The released caesium vapour would then diffuse through the chamber, which was initially kept at $\sim 80^\circ\text{C}$ ⁷. The filling (or “loading”) of the hollow-core fibres with Cs vapour is presented in Fig. 5.4. After loading the chamber for around a month, we turned the heating off, only keeping the windows at a slightly elevated temperature in order to prevent caesium accumulation. Due to the

⁶Due to incorrect wiring of the pumps and valve, the vacuum broke after a few days of loading the chamber with Cs. Luckily, after re-establishing vacuum, there was residual, unreacted metal in the Cs ampoule that eventually re-filled the chamber.

⁷Higher temperatures speed up diffusion. However, the reactivity of Cs with glass increases above $\sim 100^\circ\text{C}$, potentially increasing the chance of vacuum leaks or deterioration of the guiding properties of the HC-PCF.

large cores of kagomé-type HC-PCFs, there is a measurable trace of non-adsorbed alkali in the vapour phase even at room temperature (Fig. 5.4). Unfortunately, even in the larger $46\text{ }\mu\text{m}$ core fibre at $\sim 80^\circ\text{C}$, the ambient optical depth ($< 10^3$) was too low to perform memory experiments, since efficient ORCA operation requires optical depths⁸ on the order of 10^5 , see section 2.3.2. Fortunately, such optical depths can be generated using light-induced atomic desorption, which is the focus of the rest of this chapter.




Figure 5.4: Loading the two different core size HC-PCFs with Cs (solid lines), i.e. the change in optical depth (measured according to Appendix A. Optical depth) over time. Also shown (red dotted line) is the externally monitored temperature of the vacuum system.

5.2 Light-Induced Atomic Desorption

Light-induced atomic desorption (LIAD) has been known since the early nineties as a tool for increasing the vapour pressures of alkali atom cells [185]. For a long time it attracted relatively little interest in the atomic physics community, due in part to the popularity of bulk vapour cells, in which LIAD does not induce a substantial increase in the atomic density [194]. This is due to LIAD being a predominantly surface effect. As the surface-to-volume ratio for cylindrical cells scales as $2/r$, surface effects are not as significant in

⁸Equivalent to atomic number densities on the order of 10^{17} m^{-3} in the fibre core.

bulk cells. In $\sim \mu\text{m}$ -radius hollow-core fibres however the surface-to-volume ratio can be very large, making surface effects important. Due to caesium's extreme reactivity, a significant amount of it will be adsorbed to the core walls instead of being in the vapour phase, its clean atomic lines giving way to continuous metallic spectra. This renders adsorbed Cs unusable for implementing quantum memory protocols such as ORCA. In 2006 Ghosh et al. [195] have shown that in HC-PCFs LIAD can temporarily increase the alkali vapour pressure by several orders of magnitude. A complete, microscopic theory of LIAD has not been developed yet, though phenomenological models (analogous to the one presented in section 5.2.2) exist for particular combinations of alkalis and surfaces, including Cs on porous silica [196], Rb on porous alumina [197], Cs and Rb on silica covered with a spin-preserving coating (PDMS)[198].

5.2.1 Generating ultra-high optical depths of Cs in HC-PCFs using NIR LIAD

Seeing that even a hollow-core fibre with a large core cannot sustain the ambient optical depth of Cs vapour required to operate the ORCA memory, we have decided to re-investigate LIAD in this system. The purpose of the experiment was therefore to compare LIAD performance, using the same near-infrared (NIR) desorbing light as in [184], between two different core sizes, and optimise the process altogether, in order to increase the persistent level of LIAD-induced optical depth [95].

NIR LIAD experimental setup

A scheme of the experimental setup used to implement NIR LIAD and measure the resulting optical depths of Cs in kagomé-type HC-PCFs is shown in Fig. 5.5. The optical depth is measured by sending in a few nW⁹ of 852 nm continuous-wave (CW) probe light from an ECDL (external-cavity diode laser), scanning its frequency and measuring the absorption (see Appendix A. Optical depth). We use a combination of telescopes and achromatic lenses to couple the probe beam in and out of the HC-PCF, achieving

⁹Since the HC-PCF core sizes are $< 50\mu\text{m}$, we have to use a very low probe power in order to remain below the saturation intensity of the atomic transition ($\sim 1 - 3 \text{ mW/cm}^2$) [103]. This avoids power broadening of the probed transition making it easier to estimate the optical depth.

a coupling efficiency into the fibre's fundamental mode up to around 60% (50%) for the $26\mu\text{m}$ ($46\mu\text{m}$) core fibre¹⁰. After the HC-PCF the probe beam is reflected off a polarising beamsplitter and passes two filters: a longpass at 850 nm and a line filter at 850 nm, in order to eliminate any back-scattered 780 nm LIAD light. We measure the probe light on a slow avalanche photodiode operating in linear mode. We use a combination of precise XYZ stages holding the coupling lenses and flip mirrors to selectively couple light in and out of the $26\mu\text{m}$ and $46\mu\text{m}$ core fibres.




Figure 5.5: NIR LIAD experimental setup: HWP+QWP - half-wave and quarter-wave plates; AL - achromatic coupling lens on high-precision XYZ translation stage; HC-PCF - hollow-core photonic-crystal fibre; PBS - polarising beamsplitter; PD - photodiode; APD - avalanche photodiode.

In the NIR LIAD experiment, the desorbing beam is derived from an ECDL at 780 nm, which is mode-matched using a telescope and amplified via a tapered amplifier (TA), and sent, after an electronic shutter, into a single mode fibre, as shown in Fig. 5.6. The TA generates up to 1 W of CW light. We use a cylindrical lens and a telescope to mode-match the elliptical beam exiting the TA to the single-mode fibre, getting out up to 300 mW of NIR LIAD light at the output of the fibre. We then couple the LIAD beam counter-propagating with respect to the weak probe (combining them on a PBS¹¹), into

¹⁰After two years of operating the Cs-filled HC-PCF at room temperature, this coupling has deteriorated to around 30% for both core sizes, possibly due to the adsorbed Cs modifying the guidance properties of the HC-PCF.

¹¹We also have a photodiode on the other PBS output to record when LIAD is on.

the HC-PCF¹². The effective intra-fibre LIAD power was around 80 mW.

In order to explore the long-lived temporal dynamics of the optical depth due to LIAD in our system, we take a series of over 20,000 consecutive probe transmission scans (which correspond to three hours of data acquisition), during which we turn LIAD on and off with the electronically-controlled shutter.




Figure 5.6: Scheme of desorbing laser system for the NIR LIAD experiment. AP - anamorphic prisms; HWP - half-wave plate.

NIR LIAD experimental results

Figure 5.7 shows the change in optical depth (OD) of Cs in a $26\ \mu\text{m}$ Kagomé HC-PCF when LIAD illumination is applied at room temperature over three hours. Within the span of 10 s the OD rises by three orders of magnitude and reaches a maximum value of 1.7×10^5 , which corresponds to an average number density of atoms $7.6 \times 10^{17}\ \text{m}^{-3}$. After the sharp initial rise, the OD drops by an order of magnitude over 4 minutes. This is a significantly longer time for which such high ODs have been observed in a similar system before [184]. Furthermore, over the entire 3 hour LIAD operation the optical depth remains above 2000.

¹²We note that applying LIAD to one of the fibres did not have a measurable influence on the other.




Figure 5.7: Change in optical depth of Cs over 3 hours of NIR LIAD illumination in a $26\ \mu\text{m}$ core fibre. Insets show the behaviour of the optical depth at the beginning (left-bottom) and end (right-bottom) of LIAD illumination.

Figure 5.8 shows similar data for the $46\ \mu\text{m}$ core fibre. The initial rise in OD is slightly slower than in the $26\ \mu\text{m}$ core fibre (20 s rise time), but the final OD reached is higher, at 2.2×10^5 , after which it stays above 10^5 for up to a minute. This is comparable to the OD generated by heating a bulk cell to almost a hundred degrees (as used in chapters 3 and 4), even though the fibre system operates at room temperature, extending its operational lifetime, and in principle allowing for faster and more precise control of the optical depth (by controlling the desorbing beam power [194]). The visible mid-term fluctuations in LIAD-induced OD can be attributed to changing LIAD in-coupled powers (as is shown more explicitly in section 5.2.3), due to experimental instabilities. Switching off the LIAD beam (at $t \approx -90\text{s}$ and $t \approx -30\text{s}$ in the insets of Fig. 5.7 and 5.8, respectively) causes the OD to rapidly drop in both fibres to below initial levels.




Figure 5.8: Change in optical depth of Cs over 3 hours of NIR LIAD illumination in a $46\ \mu\text{m}$ core fibre. Insets show the behaviour of the optical depth at the beginning (left-bottom) and end (right-top) of LIAD illumination.

The LIAD-induced changes in OD appear to be governed by two processes of different timescales: a sharp initial rise and fall in OD over the course of a few minutes and a subsequent significantly slower fall in OD over the course of hours. The initial sharp peak in OD only occurs for the first run of the experiment per day and requires around a day of quiescence to regenerate to its former level, which is consistent with previous observations in similar systems [184, 192]. This is more explicitly visible in Fig. 5.9, which shows the change in optical depth when LIAD illumination is applied in 50 s pulses separated by 10 s. Here, only the first LIAD pulse of the series generates the ultrahigh optical depth, with the consecutive pulses yielding a slowly decaying “steady state” OD. Additionally, the OD rise time is longer for these consecutive pulses, suggesting a different mechanism of desorption. These different desorbing mechanisms with their specific temporal dynamics

are explained in the next section, in which I attempt to model LIAD.




Figure 5.9: Change in optical depth of Cs (blue) during pulsed NIR LIAD illumination (orange) in the $26\text{ }\mu\text{m}$ core fibre.

5.2.2 Modelling the NIR LIAD experiment

Most authors seem to agree on a theory postulating the existence of adsorbed caesium in two forms - metallic layers and nanoclusters [199, 192] (some authors also stress the importance of the caesium atoms absorbed into the coating or wall medium [200, 197]). The metallic layer interacts with the desorbing light due to a non-resonant, photoelectric-type mechanism, while the nanoclusters evaporate due to surface plasmon resonance-enhanced heating. Hoheisel et al. suggest that both desorption mechanisms have the same underlying principle and the metallic layer simply does not experience the strong optical field enhancement due to the plasmon resonances [201].

According to Bhagwat et al. the metallic layer contains a smaller number of desorbable alkali atoms, but regenerates very quickly, producing a steady state level of OD, as the atoms continuously desorb and re-adsorb on the fibre core walls [192]. The nanoclusters on the other hand, are very efficiently desorbed, generating large ODs on very short

timescales, but re-adsorb as layers, and take hours of quiescence to reform on wall surfaces as nanoclusters. Most of the experiments on LIAD were conducted at moderate light intensities however and it has been suggested that at high intensities thermal effects can become important [199, 192].

The present knowledge on LIAD in fibre geometries does not give a clear answer on how to use it efficiently for obtaining persistently high values of the optical depth of Cs in our HC-PCF setup. In fact, it suggests that the LIAD yield is very sensitive to the particular experimental conditions, including the desorbing light frequency, intensity, surface chemistry, geometry and temperature.

I have thus decided to describe LIAD in our system using a phenomenological model (shown in Fig. 5.10) based on rate equations. I originally assumed Cs could exist in three forms in the experiment, i.e. be adsorbed on the fibre core walls as a metallic layer or nanoclusters, or exist in the vapour phase. A strong off-resonant beam then desorbs the layers and nanoclusters via two mechanisms of different desorption rates. Cs in the vapour phase can then re-adsorb to the core walls, but only via formation of metallic layers, which then reform into nanoclusters at a certain slow rate.




Figure 5.10: Postulated model of LIAD. The four forms of Cs are shown: (a) vapour, (b) metallic layers, (c) nanoclusters, (d) sink, along with the processes of transforming between them.

However, considering only the desorption of nanoclusters and layers, and their refor-

mation/readsorption, could not explain the slow fall in OD observed in our experiment. In order for our model of LIAD to mimic the data correctly, we had to include another process. This process decreases the amount of Cs in the whole system under LIAD illumination, but regenerates it during the period of darkness. In coated bulk cells, the process responsible for similar long-term depletion of desorbable alkali was identified to be its interaction with the coating itself, including light-stimulated absorption into the coating [200]. A similar depletion effect was also observed in LIAD from porous alumina, where the pores acted as a sink for the atoms [197], and in vapour cells with closable stems [194]. We speculate that in glass hollow-core fibres the mechanism responsible for long-term depletion of alkali vapour during LIAD illumination is the escape of the vapour out of the fibre ends, due to a strong, LIAD-induced vapor pressure differential on the boundary between the fibre and its surroundings.

We thus arrive at the following rate equations describing our system:

$$\begin{aligned}
\partial_t N_v &= R_1^{\text{LIAD}} P_{\text{LIAD}} N_l + R_c^{\text{LIAD}} P_{\text{LIAD}} N_c \\
&\quad - R^{\text{ads}}(N_v - \mu_l N_l) - R^{\text{esc}}(N_v - \mu_s N_s) \\
\partial_t N_l &= -R_1^{\text{LIAD}} P_{\text{LIAD}} N_l + R^{\text{ads}}(N_v - \mu_l N_l) \\
\partial_t N_c &= -R_c^{\text{LIAD}} P_{\text{LIAD}} N_c \\
\partial_t N_s &= R^{\text{esc}}(N_v - \mu_s N_s)
\end{aligned} \tag{5.1}$$

where N_v, N_l, N_c are the fractions of atoms in the vapour, layer and nanocluster phase, respectively. The first two terms on the right-hand side of the first equation describe LIAD, i.e. the light-induced desorption of metallic layers and nanoclusters. The desorption rates are R_1^{LIAD} and R_c^{LIAD} for the layers and nanoclusters respectively, with nanocluster desorption being much more efficient ($R_c^{\text{LIAD}} \gg R_1^{\text{LIAD}}$). P_{LIAD} is the LIAD beam power. The next two terms describe the equilibrium processes of alkali adsorption as layers and its escape out of the fibre as vapour. R^{ads} and μ_l are the adsorption rate (fast) and equilibrium ratio of atoms in the vapour and layer phases. N_s is the fraction of atoms in the sink, i.e. outside the fibre, and R^{esc} and μ_s are the atomic escape rate from the fibre (slow) and the equilibrium ratio of atoms on the outside and inside the fibre (in the vapour phase). It is worth noting that we have omitted the clusterisation

phenomenon, shown in Fig. 5.10, from the equations. Its effect on the results presented in Fig. 5.11 and 5.12 was found to be negligible. This is because clusterisation occurs on timescales much longer than the data acquisition time, the nanoclusters requiring a day of quiescence to reform compared to the few hours of LIAD illumination.




Figure 5.11: Comparison of LIAD-induced change in optical depth over time for data from the 26 μm core fibre (top) and our rate-equation model (bottom). Insets show the behaviour of the optical depth at the beginning (left) and end (right) of LIAD illumination.

Table 5.1: **Rate coefficients from Eq. 5.1 used to generate the model shown in Fig. 5.11.**

R_c^{LIAD}	R_l^{LIAD}	R^{ads}	R^{esc}
6	0.06	15×10^{-2}	16×10^{-7}

Figure 5.11 shows a comparison between experimental data (top) and the numerical solution of the model described in section 5.2 (bottom). The model was fed with the

measured LIAD power in the fibre, and the rate coefficients from Eq. 5.1 were found manually (shown in Table 5.1), in order to get the closest similarity to data. Indeed, the model accurately describes the qualitative features of the data, with the relations between the magnitudes of the fitted rate coefficients being as expected, e.g. the nanocluster desorption rate being much higher than the layer desorption, i.e. $R_c^{\text{LIAD}} \gg R_l^{\text{LIAD}}$; and the rate of escape of atoms into the sink being much smaller than the adsorption rate ($R^{\text{ads}} \gg R^{\text{esc}}$), having a visible effect only during LIAD when the density of vaporous atoms in the fibre is large. Furthermore, using our model we can confirm that this reversible escape of atoms into the sink during prolonged LIAD illumination, which physically could be the escape of the atomic vapour into the chamber out of the fibre ends, is the main process limiting the long-term OD.




Figure 5.12: Comparison of LIAD-induced (orange) change in optical depth (blue) over time for data from the 26 μm core fibre (top) and our rate-equation model (bottom).

We also model LIAD in the pulsed illumination regime (50 s pulses separated by 10 s), as shown in Fig. 5.12. Similarly, we find the rate coefficients from Eq. 5.1 manually to get the closest similarity to data. Again, the model reproduces the qualitative features of

Table 5.2: **Rate coefficients from Eq. 5.1 used to generate the model shown in Fig. 5.12.**

R_c^{LIAD}	R_l^{LIAD}	R^{ads}	R^{esc}
1	0.2	2×10^{-2}	9×10^{-7}

the data quite well, such as the generated peak in the optical depth (due to nanocluster desorption) unique for the first LIAD pulse of the series, and the slow fall in long term maximum OD generated per pulse, due to the atoms escaping into the sink. The rate equation coefficients (listed in Table 5.2) are also not very far off from the ones in Table 5.1, and maintain similar relations between the different coefficients, as expected from two experiments carried out in the same system.

Our rudimentary numerical modelling of the experimental data confirms our intuitions about the processes occurring during LIAD in an alkali-filled hollow-core fibre. While nanocluster desorption leads to very high transient optical depths, layer desorption provides a moderate, but persistent level of optical depth, which is desirable for single-photon experiments requiring long data acquisition rates (as described in chapter 4). Furthermore, using our model we were able to identify the loss of atoms, possibly due to their escape from the fibre during LIAD illumination, as the main process responsible for the long-term fall in OD. A possible way to prevent the long-term drop in OD is to seal the alkali-filled HC-PCF ends to prevent the atoms inside from escaping. The duration of the transient ultrahigh ODs due to nanocluster desorption would still be limited by the long reformation time of the nanoclusters. It would then be of great value to increase the efficiency of layer desorption, in order obtain persistently high optical depths. It has been suggested that layer desorption becomes more efficient at lower wavelengths [199], which we consequently endeavoured to investigate in our system.

5.2.3 Generating even higher optical depths using green LIAD

Studies in bulk glass cells have shown that while nanocluster LIAD is maximised for specific frequency bands in the NIR, due to plasmon resonances of the nanoclusters, metallic layer desorption efficiency rises exponentially with the frequency of the desorbing light [199]. This would naively suggest that using lower wavelength light could improve the generation of persistently high optical depths in alkali-filled hollow-core fibres by improving

the LIAD yield of the “persistent” desorption process, as we understood layer desorption in section 5.2.2. Of course, without changing the equilibrium amount of caesium in the metallic layers compared to the other forms¹³, improving the desorption efficiency of the layers will simply lead to faster depletion of that amount. Nonetheless, we have decided to investigate LIAD with considerably higher frequency green light, since this has never been attempted in alkali-filled hollow-core fibres before¹⁴, and could yield potential improvements in LIAD performance¹⁵.

Green LIAD experimental setup

In order to improve layer desorption in Cs-filled HC-PCFs, we have repeated the experiment described in section 5.2.1 with green desorbing light, instead of NIR. In this case, the desorbing beam was derived from a 5 W continuous-wave “Millennia” laser¹⁶. We couple the 532 nm light from the Millennia into a single-mode fibre (SMF)¹⁷ and send it to a slightly modified version of the setup shown in section 5.2.1, as shown in Fig. 5.5. In this experiment, we separate the weak 852 nm probe from the strong counter-propagating 532 nm LIAD beam using a dichroic mirror. We also pick off some of the LIAD beam after the HC-PCF, using a glass plate, in order to continuously monitor the in-coupled LIAD power. It was not immediately clear from the fibre transmission spectrum shown in Fig. 5.1 (section 5.1.1) that our HC-PCFs would couple and guide efficiently 532 nm light. However, we measured the total transmission of the 26 μm ¹⁸ core HC-PCF at 532 nm (including coupling efficiency, using the same optics as for the NIR LIAD) to be around 28%, giving us up to around 600 mW of 532 nm light inside the fibre (with the Millennia power output set to 4 W). The measurement of the LIAD-induced change in OD over time was performed analogously as described in section 5.2.1, with the difference being

¹³This was shown in bulk cells to be possible by two-color illumination, which promoted the formation of nanoclusters or layers [199].

¹⁴The reason being that most experiments investigating LIAD in hollow-core fibres used photonic bandgap fibres with narrow bandgaps, precluding the possibility of efficiently guiding very far off-resonant light [202].

¹⁵Especially considering the higher amount of power we had at our disposal in the green, see next section.

¹⁶The Millennia laser was once used to pump an old Ti:Sapph laser.

¹⁷Although we managed without a dedicated optical isolator for the Millennia, it was crucial to use angled fibres in order to get rid of back-reflections feeding into the laser at high powers and causing laser power fluctuations.

¹⁸The measurements described in this section were performed solely in the 26 μm core fibre, since both fibres should exhibit similar behaviour, as we have seen from the results in section 5.2.1.

the current monitoring of the in-coupled LIAD power. This allowed us to improve the coupling of the Millennia SMF and control “in-situ” the coupling of the desorbing beam into the HC-PCF, both of which would fluctuate significantly (as we will see in the next section), if the Millennia output power was set to ≥ 500 mW.




Figure 5.13: Green LIAD experimental setup: HWP+QWP - half-wave and quarter-wave plates; DM - dichroic mirror; HC-PCF - $26 \mu\text{m}$ hollow-core photonic-crystal fibre; PD - photodiode; APD - avalanche photodiode.

Green LIAD results

We first implement the 532 nm LIAD at a low in-coupled desorbing beam power of ~ 25 mW¹⁹. The LIAD-induced change in optical depth over time (blue line), along with the change in the LIAD in-coupled power (orange line) is shown in Fig. 5.14. Noticeably, the in-coupled LIAD power deteriorates over the course of the 1 hour long measurement quite significantly ($\sim 50\%$) even though it was peaked up several times throughout the measurement (the visible step increases in the orange line at ~ 2 min. and ~ 18 min.²⁰). As with the NIR LIAD experiment, we see a sharp rise in optical depth due to green LIAD, followed by a rapid fall and then a slow decrease in the optical depth over time. It is difficult to directly compare the results in this section to the ones presented in section 5.2.1, since the two experiments have been performed over a year apart, during which

¹⁹The Millennia power output set to 1 W.

²⁰The sharp dips in the orange line are due to attempts to peak up the system that would briefly deteriorate coupling of the light.

the amount of Cs inside the HC-PCF could have changed²¹. However, we can make the claim that we did not see the increase in the “steady state” OD (even at low power) due to LIAD at 532 nm, which we have initially hoped for, and after the initial peak in OD, the optical depth stabilised to on the order of 10^3 , similarly to the NIR experiment.




Figure 5.14: Change in optical depth of Cs (blue line) over 1 hour of green LIAD illumination in a $26\ \mu\text{m}$ core fibre. The orange line depicts the fibre in-coupled LIAD beam power.

In order then to see whether we could at least improve the peak OD by using more desorbing beam power (which was available to us in the green), we repeat the measurement with up to $\sim 600\ \text{mW}$ of in-coupled desorbing beam power²², the results of which are shown in Fig. 5.15. We break our previous OD record from section 5.2.1 by a factor of 5, reaching a maximum optical depth of around 1.1×10^6 . It is difficult however to separate the contributions of longer fibre loading time, higher desorbing beam power, and LIAD

²¹The NIR desorbing laser system was inoperational during the green LIAD measurement, and so it was impossible to measure and compare LIAD using the two wavelengths directly.

²²The Millennia power output set to 4 W.

wavelength. The rise time of the peak OD is 10 s, after which it drops by roughly two orders of magnitude over the course of a minute, which is noticeably faster than in the case of NIR LIAD.




Figure 5.15: Change in optical depth of Cs (blue line) over 1 hour of high-power green LIAD illumination in a $26\ \mu\text{m}$ core fibre. The orange line depicts the fibre in-coupled LIAD beam power. Insets show close ups of certain sections of the data discussed in the main text.

At some points during (roughly 100 s and 1500 s into) the high power measurement we have consciously deteriorated the coupling of the desorbing beam into the HC-PCF. This caused a substantial increase in the optical depth generated in the fibre, which followed the regular pattern of pulsed LIAD (see Fig. 5.9 in section 5.2.1), i.e. with the first “decoupling” — a peak in OD, the maximum of which is on the order of the “coupled” one, i.e. $\sim 10^6$, followed by a slow fall — and then for the second — a much weaker rise in OD, followed by a slow fall. We postulate the decoupling leading to excitation of higher order modes in the HC-PCF by the desorbing beam, which thus interacts with different parts of the fibre core walls, desorbing a priorly un-desorbed region of the core. Interestingly, this region seems to contain a significant amount of desorbable atoms, as

it sustains optical depths above 10^4 for almost 20 minutes. Re-coupling the HC-PCF (about 900 s into the measurement), the optical depth would actually drop, and even at a significant amount of in-coupled desorbing beam power, in 10 minutes the OD would fall from 2000 to around 700. This confirms the suspicion that depending on which fibre mode the LIAD beam couples into, different pockets of adsorbed Cs are desorbed and, once depleted, cannot be immediately desorbed again. Thus by modulating the coupling of the desorbing beam throughout LIAD, it ought to be possible to generate more Cs vapour. In the decoupled regime however, LIAD becomes more sensitive to experimental instabilities.

In conclusion, even though we expected metallic layer desorption (responsible for the “steady-state” LIAD) to be much more efficient in this experiment, compared to the NIR one (due to the change in wavelength and improvement in desorbing beam power), we have not been able to significantly improve the long-term performance of LIAD. As expected from literature [198, 199], the dynamics of LIAD have become faster by switching to a lower wavelength, but this has not lead to higher steady-state values of the optical depth. We have however identified the excitation of specific transverse modes of the HC-PCF by the desorbing beam as a potential avenue which could lead to an improvement in LIAD performance.

In summary, I have investigated kagomé-type hollow-core photonic crystal fibres filled with warm Cs vapour as a potential platform for integrating the ORCA quantum memory. While we demonstrated ultra-high transient optical depths of Cs in this system using light-induced atomic desorption, we were not able to show a persistent level of optical depth, sufficient to perform time-extensive quantum memory experiments (as described in chapter 4).

6

Conclusion

Photonic quantum networks are a very versatile platform for quantum research and engineering, as they can be applied to quantum communication, metrology and computation [1]. The quantum properties of light are not degraded by noise in ambient conditions, and photonics can support large bandwidths enabling fast operations and a large information capacity [5]. However, the current paradigm for photonic quantum information processing, which relies on single photons, linear optics and measurement, is not scalable [7], unless it includes quantum memories, i.e. devices that can store and release photons on-demand [203].

Historically, the goal of sharing entanglement over global networks motivated the development of “quantum repeaters”, atomic quantum memories with sufficiently long memory lifetimes to synchronise entangling operations over lossy, high latency networks [38]. In contrast to long-distance communications networks, local networks for quantum

information processing exhibit low latency. The key figures of merit for memories in local quantum networks are therefore the efficiency and the time-bandwidth product, i.e. the number of clock-cycles over which a memory allows synchronisation [55]. Regardless of the unique applicability of each memory technology, there is an additional overarching requirement for quantum operation in that the memory must be noise free.

While low noise, large time-bandwidth products, and operation in ambient conditions have been shown separately using various memory protocols, to-date no atom-based memory satisfies all these desiderata simultaneously. Here I introduced and demonstrated a new memory protocol — the off-resonant cascaded absorption (ORCA) memory — that is noise-free in ambient conditions, technically simple and thus potentially integratable, and can support sufficient time-bandwidth products for local quantum network technologies.

6.1 Summary of thesis

First, in chapter 2, I presented a theoretical investigation and characterisation of the ORCA memory, deriving the equations of motion that govern it, solving them in certain parameter regimes analytically, and then numerically in a broad parameter range. I found that the memory should be capable of reaching the performance levels required for certain applications, such as temporal multiplexing. I then implemented the ORCA memory in warm caesium vapour using $\sim$ GHz-bandwidth near-infrared pulses, which is the topic of chapter 3. I have characterised the memory using single-photon level coherent states, measuring a time-bandwidth product of ~ 4 , and achieving memory efficiencies of up to $\sim 17\%$, in line with our theoretical predictions for the current experimental parameters. Most importantly though I measured no noise being generated by the memory.

It is important to note that it is insufficient to predict guaranteed quantum operation by only measuring the mean of the noise, because even a very small average amount of noise can impair quantum signatures if the variance of the noise is large i.e thermal [16]. I have therefore interfaced the ORCA memory with a heralded single-photon source based on parametric downconversion in a ppKTP waveguide. The photon number statistics of this quantum state are a very sensitive probe of noise. In chapter 4, I consequently verified that the ORCA memory is indeed noise-free by measuring and comparing the normalised

Glauber correlation functions, specifically the cross-correlation $g^{(1,1)}$ and heralded auto-correlation $g_h^{(2)}$, of the input and output light. I found that, within the measurement accuracy, both functions were preserved, measuring a record low $g_h^{(2)} = 0.03(1)$ upon retrieval. I have thus demonstrated for the first time an inherently noise-free broadband on-demand atomic quantum memory, let alone in a technically simple implementation operating in ambient conditions [146]. From a more fundamental standpoint, I have showed that it is possible to strongly interact a single photon with a collection of ~ 10 billion hot atoms and a classical field composed of $\sim$ billion other photons, and still fully preserve the photon's quantum properties.

Finally, I made some headway into integrating the ORCA memory by investigating an alkali-filled hollow-core fibre architecture, as described in chapter 5. While originally designed for the Raman memory [15], the large-core kagomé-structured fibres are very well suited to the ORCA memory. First of all, the latter is inherently noise-free and thus does not require further engineering to reduce the noise (such as a cavity [140]). It does however require two different wavelengths to operate (and the fibres are broadband), and is limited in lifetime to hundreds of nanoseconds in warm alkali vapours (which is also the transit time of the atoms through the fibre core). Finally, ORCA does not require optical pumping, which has proved a challenge in the fibre system¹. I have thus rebuilt an ultra-high vacuum system, which housed two fibres of different core sizes, and passively loaded it with warm caesium vapour. I found that the ambient optical depth of caesium inside both fibres was limited by the adsorption of the alkali to the fibre core walls. Using near-infrared (NIR) and green light-induced atomic desorption (LIAD), I was able to generate record high transient optical depths of the alkali inside the fibres, up to $\sim 10^5$ and $\sim 10^6$ with the NIR and green desorbing light, respectively. In order to better understand the limitations of LIAD in our system, I developed a rate-equation model of it, and identified the loss of atoms, possibly due to their escape from the fibres during LIAD illumination, as the main process responsible for the long-term fall in atomic number density [95].

It is important to realise that this work is the culmination of over ten years of quantum memory research in the Ultrafast Quantum Optics group at the University of Oxford [134].

¹I have found, re-implementing the Raman memory in the Cs-filled fibres, that optical pumping in this system is not very efficient, and also leads to dephasing of the memory.

Metric	State-of-the-art	Demonstrated	Expected
Efficiency	87% [204]	17%	60%
Lifetime	40 s [125]	5 ns	100 ns
Bandwidth	2 THz [133]	1 GHz	> 1 GHz
Time-bandwidth product	$\sim 10^5$ [125, 11]	~ 5	> 100
Noise level (μ_1)	3×10^{-3} [126]	$< 10^{-5}$	0
Fidelity	98% [105]	17%	60%

Table 6.1: Comparison of the performance of state-of-the-art memories and the current ORCA demonstration. Fidelity is calculated in the photon number basis³. Also shown is the expected possible performance of ORCA.

For a long time, the main research goal was to develop an efficient room-temperature quantum memory with high time-bandwidth product based on the Raman memory protocol [17]. However, in the process of reaching that goal, noise was identified as the critical issue faced by not only the Raman [16], but many memory protocols, and preventing quantum memories from being applied in a quantum information processing scenario. ORCA finally solves the problem of noise [146], opening the way to using quantum memories to build large photonic quantum networks, with broad applications in quantum simulation, computation, and sensing.

The performance of the present proof-of-concept implementation of the ORCA memory, in terms of the figures of merit described in section 1.1, is summarised in Table 6.1. Comparing this to the current state-of-the-art (also shown in Table 6.1), the ORCA memory is currently outperformed in all but noise. However, as detailed in section 6.2.2, its performance in most metrics can still be improved (last column in Table 6.1), illustrating the great potential of this new memory protocol.

6.2 Outlook

While representing a significant conceptual step, the demonstrated performance, at least in terms of memory lifetime (section 3.2) and efficiency (section 3.3), of our ORCA implementation in warm caesium vapour is appreciably insufficient for practical quantum information processing applications (section 1.2.4). Our theory however predicts that the

³In the ORCA case, I calculate the fidelity for storing single photons, i.e. the overlap between the desired single-photon Fock state $|1\rangle$ and the produced $|\Psi_{mem}\rangle \approx \sqrt{\eta}|1\rangle + \sqrt{1-\eta}|0\rangle$ (this assumes there is zero added noise, as evidenced by our measurements).

ORCA memory performance could be improved (section 2.3.2). Furthermore, there are still a few properties of the ORCA protocol that we do not fully understand theoretically or have not verified experimentally. Finally, full integration of the ORCA memory carries its own challenges. Here I list the next few potential research areas of interest which directly follow from the work presented in this thesis.

6.2.1 Investigating the modal capacity of ORCA

As has been mentioned in chapter 2, different applications require quantum memories to be either single-mode (e.g. for spectral-temporal mode filtering [62]) or multimode (for quantum communications [115]). While the modal properties of e.g. the Raman [136] or GEM [106] memories are well understood, the ORCA memory has not yet undergone the same level of theoretical scrutiny. In chapter 2, I briefly discussed how the modal properties of ORCA should mimic the Raman memory in the short interaction length limit, making it single-mode for low coupling strength, and multimode at strong coupling [136]. Though, as mentioned in section 1.3.2, single-mode operation can be combined with high memory efficiency using a cavity [140]. In the same chapter, I mentioned that in the long cell limit, when the counterpropagating signal and control pulses completely pass each other, the ORCA memory should be multimode in time, with different time bins being stored at different positions in the cell (similarly to how different frequency bins are stored in GEM [98])⁴. Of course, the latter result was derived under the assumption of the absence of the AC Stark effect, which limits the memory efficiency in realistic scenarios [140].

It would therefore be interesting to include the effect of the AC Stark shift in the long cell regime, and see whether the good multimode properties of ORCA survive. Furthermore, it would be insightful to investigate the intermediary cell length regime analytically, and the transition from a multimode to single-mode memory. The same transition could possibly be observed by moving the point of maximum temporal pulse overlap of the signal and control fields between the edge of the cell and the middle, as this would effectively

⁴It would be interesting to investigate whether this property of ORCA could be exploited to create a linear array of atomic excitations, which (once transferred to a Rydberg state) could be used for quantum simulation of e.g. interacting spin chains [205].

change the memory interaction length.

6.2.2 Improving the ORCA memory performance

While the noise properties of ORCA are superior to any other atomic quantum memory system to-date, the demonstrated memory lifetimes and efficiencies do not reach the same level as some of the other memory protocols (see section 1.3), or even a level sufficient for multiplexing of heralded photon sources (section 1.2.4) [55]. Fortunately, our theoretical analysis suggests that the memory efficiency could be improved by a factor of ~ 4 by optimising the vapour cell temperature, detuning, signal/control temporal overlap and control pulse energies (section 2.3.2). Furthermore, our results seem to suggest that this experimentally optimal efficiency is limited by the AC Stark effect, as was found to be the case in experimental implementations of the Raman memory⁵ [17, 140]. As with the latter, a potential solution to this is appropriate shaping of the control field pulse [206]. A preliminary mode analysis of the ORCA memory (using a numerical Green's function approach [23]), has shown that indeed by appropriately matching the signal and control pulse shapes, efficiencies over 80% could be reached. Combining this with an optimisation of the losses in the setup, it is feasible that end-to-end device efficiencies of $\sim 80\%$ are within reach, close to the current state-of-the-art [41].

In the case of long-lived, narrowband memories, the control field can be shaped using acousto-optic (tens of megahertz) [116] or electro-optic (up to a gigahertz) [207] modulators. Dispersion based techniques are also well-known for fs-long pulses [208]. The GHz-THz bandwidth regime is thus a very tricky one to shape pulses in. It would therefore be beneficial to increase the operating bandwidth of the ORCA memory⁶, especially since it should not be limited by the ground-state splitting in alkalis, as it is in Λ -based protocols (see section 1.3.2). It would however be necessary to first theoretically investigate the effect that increasing the memory bandwidth would have on memory efficiency and lifetime (via hyperfine dephasing, see brief discussion in section 3.2.2), which would

⁵There though the efficiency was also effectively limited by noise introduced by the gain due to four-wave mixing [100], which is absent in ORCA.

⁶Increasing the memory bandwidth could also yield other advantages for some of the applications outlined in section 1.2, however for multiplexing there is no direct reason to increase the bandwidth of the stored light beyond the bandwidth of available feed-forward electronics ($\sim$ GHz) [18].

require including both ground states, and the whole excited and storage state hyperfine manifolds in the simulation.

Improving the memory lifetime is another important research goal. We have found that in the current implementation in warm caesium vapour the ORCA memory lifetime was limited by motion-induced dephasing and hyperfine dephasing (see section 3.2) to a few nanoseconds. Longer lifetimes would need to be demonstrated for the memory to be useful for quantum information processing applications, let alone quantum communications. Hyperfine dephasing, as explained in section 3.2.2, could be eliminated by appropriate preparation of the atomic ensemble and choice of signal and control polarisations, such that only one transition pathway to a specific hyperfine storage state is possible⁷ [18]. Reducing motion-induced dephasing is trickier. One obvious solution is to cool the atomic ensemble, using laser cooling methods [145]. Another would be to prepare the ensemble in one of the longitudinal velocity classes via velocity-selective pumping, in order to reduce Doppler broadening (which is behind motion-induced dephasing) [209]. In both cases however, it is not immediately clear that sufficient optical depth could be achieved. An altogether different method would be to transfer the doubly-excited coherence back to a meta-stable ground state after storage, and then re-map it to a doubly-excited coherence before attempting retrieval. The potential noise sources in this scheme would have to be investigated however. There may also exist ways to rephase motion-induced dephasing [209].

Another approach would be to investigate other storage materials. In section 3.2.2, we noted that by implementing ORCA in warm rubidium vapour, where motion-induced dephasing is reduced owing to a smaller signal/control wavelength mismatch, a memory lifetime of ~ 100 ns could in principle be reached. In fact, our collaborators at the Weizmann Institute of Science in Israel have recently demonstrated a lifetime of around 85 ns in a fast ladder memory (FLAME) — equivalent to ORCA when far-detuned — using classical pulses in rubidium vapour, confirming our theoretical expectations [18]. Of course, solid-state materials such as rare-earth ion-doped crystals would exhibit no motion-induced dephasing, making them an interesting storage medium to investigate in

⁷Using standard optical pumping methods for ensemble state preparation could however limit the ORCA memory's operating speeds.

the context of ORCA (and there are known experimental techniques to operate within the large inhomogenous broadening present in those systems, as discussed in section 1.3.1).

Finally, we have not investigated the memory fidelity in the current work, which is an important figure of merit for many of the applications outlined in section 1.2. Because of the low memory efficiency, we anticipate the Fock-state fidelity to be comparably low (even though ORCA does preserve the photon number statistics.). As discussed in Appendix D. Polarisation, because of the complex polarisation selection rules present when storing light in a whole manifold of hyperfine states, we expect the polarisation state fidelity of the current ORCA implementation to be low as well. Finally, our simulations show that the temporal shape of the stored signal changes quite significantly at high coupling strengths (see section 2.3.2), which would also deteriorate the temporal mode fidelity, and potentially indistinguishability, of stored photons. Further work thus has to be done on improving the fidelity of the ORCA memory by improving efficiency, investigating the polarisation properties of the memory, and appropriate shaping of the control pulse.

6.2.3 Verifying ORCA quantum operation

In the current work, we have verified that ORCA can be applied for the storage of light, and confirmed that it is a noise-free light-matter interface by demonstrating preservation of photon number statistics. We have not however explicitly verified the coherent operation of the ORCA memory⁸, which could be done by interfering a retrieved coherent state with a non-stored time-delayed copy [17]. Furthermore, it would be interesting to attempt more complex quantum operations using ORCA, such as storing half of an entangled state [210]. As has been mentioned in the previous section, the polarisation fidelity of the current implementation of the ORCA memory would probably be too low to demonstrate storage of polarisation entanglement. However, storage of time bin entanglement could be attempted [85]. This would involve storing one photon of the Bell-like state $|\Psi\rangle = |ee\rangle + |ll\rangle$, which is produced when a probabilistic (e.g. PDC) photon source

⁸We do however have strong indirect evidence of coherent operation, such as similarity to our theoretical model, which assumes the creation of atomic coherence, as well as the fact that we observe efficient retrieval, which would be impossible if the control field only interacted with a single atomic excitation with no collective coherence. Furthermore, the Raman memory was demonstrated to be coherent [17], and the ORCA memory, being an off-resonant scheme as well, should share many of the former's properties.

is pumped by two time-separated (early and late) pulses from a coherent light source [211]. The source then produces a coherent superposition of the pair of photons being generated by the early pulse ($|ee\rangle$) or the late one ($|ll\rangle$). One would then attempt storage and retrieval of one of these photons, and measure the final state of both photons using a pair of Franson interferometers, in the $|e\rangle\langle e|/|l\rangle\langle l|$ basis of both photonic modes [85].

Another interesting experiment made possible by storing single photons in ORCA would be to use the off-resonant memory as a light-matter beamsplitter to demonstrate Hong-Ou-Mandel Interference (HOMI) [61]. HOMI is a purely quantum phenomenon, which occurs when two indistinguishable single photons⁹ impinge on a beamsplitter [212], i.e. at the input to the beamsplitter we have a state $|\Psi\rangle = |1\rangle_a|1\rangle_b$ ¹⁰. Due to quantum interference, the photons at the output of the beamsplitter will be bunched, both leaving in only one of the two output modes, i.e. ideally forming a state $|\Psi\rangle = (|2\rangle_c|0\rangle_d - |0\rangle_c|2\rangle_d)/\sqrt{2}$ (a two-photon NOON state). Placing detectors at both outputs of the beamsplitter, one should therefore observe no coincidences between the detector clicks. In the case of a light-matter beamsplitter, one would like to observe quantum interference between the stored single matter excitation and a single photon. It is non-trivial however to directly measure a matter excitation [84]. Hence, memory-mediated HOMI was observed in a bulk diamond Raman memory using a quasi-photon-number-resolving detector on the memory output, in order to verify the prevalence of the two-photon component of the optical mode, indicative of quantum interference [61]. With a $\sim 100\%$ efficient memory one should in principle be capable of re-mapping the HOMI-generated two-photon component of the matter excitation back to an optical field at a different time, and directly measuring (the absence of) coincidences between the initial and read-out time bins.

6.2.4 Improving ORCA integration

For practical applications, one would desire an integrated and energy-efficient quantum memory, which could be easily interfaced with existing quantum photonic circuits. This was the original motivation behind implementing the Raman memory in an alkali-filled

⁹In fact, HOMI is often used as a measure of photon indistinguishability.

¹⁰We label the two input modes of the beamsplitter a and b , and output modes c and d .

hollow-core fibre architecture¹¹, where it was shown that the fibre-based memory required 200 times less control pulse energy than one based in a bulk vapour cell (thanks to the higher Rabi frequency with the tighter focusing) [184]. To some degree these benefits of the fibre architecture should transfer to the ORCA memory, with others specific to the latter (as has been mentioned in section 6.1). In chapter 5 however, we have identified a significant challenge faced by the alkali-filled fibre platform, i.e. the difficulty (even when using LIAD) in sustaining persistent high optical depths, as required if one wishes to conduct time-extensive experiments such as photon counting (see section 4.2.1), or generally use the fibre-based memory for practical applications.

Persistent high optical depths have not been observed to-date in warm alkali-filled fibres¹². As we have seen in section 5.1.2, even operating the fibre system at close to 100°C does not yield sufficient ambient optical depth, and one is required to use LIAD to generate the necessary atomic number density. We have found that both short wavelength and long wavelength LIAD only leads to ultrahigh transient optical depths, with the long-term increase being much more moderate. In bulk alkali vapour cells, persistent increased optical depths were achieved by modulating the desorbing beam power [214]. A similar technique could be investigated in hollow-core fibres. In general, it would be enlightening to study the dependence of LIAD performance on desorbing beam power and investigate pulsed LIAD operation to a greater extent [202]. It would also be interesting to investigate LIAD when the desorbing beam is coupled to higher order modes of the hollow-core fibre (which we have briefly discussed in section 5.2.3). Since LIAD is a surface effect and higher order modes have a greater overlap with the fibre core surface, one could imagine improving LIAD performance by using higher order modes.

The level of achievable persistent optical depths will ultimately be limited by the available amount of caesium in the fibre core. Using our rate equation model of LIAD (section 5.2.2), we have seen that during prolonged LIAD illumination the amount of caesium in the core decreases. We postulated that this is because caesium vapour escapes out of the fibre ends into the vacuum chamber, due to the strong LIAD-induced pressure

¹¹However, the noise present in the fibre-based Raman memory prohibited its quantum operation [15].

¹²Quasi-persistent large optical depths have been achieved with cold atomic clouds dropped into hollow-core fibres [213].

differential at the fibre ends [95]. Sealing the fibre ends would therefore potentially improve persistent LIAD. While there has been progress in splicing large-core hollow-core photonic crystal fibres to regular single-mode solid-core fibres [215] (which has the added benefit of potentially making the alkali-filled fibre memory plug-and-play), this would be challenging to do with an alkali-filled fibre. Due to the alkali's chemical reactivity, the tapering and splicing operation would have to be performed in ultra-high vacuum. A different option would be to use a thermally decomposable compound of caesium that could be somehow inserted into the fibre under ambient conditions, enclosed via splicing to single-mode fibre, and then heated to release the caesium. This would potentially change the surface chemistry of the fibre core though, changing LIAD behaviour. There has also been progress in creating fibre-based microcells using various encapsulation methods [216, 217], which in principle could be performed in vacuum. Finally, there have been attempts to couple warm vapours evanescently to nanophotonic structures [218]. On the other hand, implementing the ORCA memory in a solid-state system, such as rare-earth ion-doped crystals, opens a whole other avenue of integration, where waveguides [219] or even nanophotonic structures [26] can be directly etched into the storage material.

Finally, room for improvement could be found in the process of loading the fibre with caesium, as currently it takes a few months to fill it sufficiently, significantly lowering the prototyping rate of these types of experiments. One could think of using a pressure differential (e.g. by introducing a buffer gas in part of the chamber) to speed up the loading of the fibre core. The process of filling the fibre with caesium and the depth of filling are actually not well understood, and it may be that we only observe caesium in small regions at the fibre ends [220]. The degree of fibre filling is especially crucial for investigating the multimode regime of the ORCA memory (as described in section 6.2.1), since the fibre system would be ideal for implementing the memory in the long cell limit. It would thus be interesting to develop a method to measure the longitudinally-resolved optical depth of caesium, possibly using an external localised magnetic field and looking at the Faraday rotation of light [221] at different positions in the fibre, or using a highly-localised short-pulse ORCA interaction.

References

- [1] H. J. Kimble. The quantum internet. *Nature*, 453(7198):1023–1030, 2008. [1](#), [118](#)
- [2] A. O. Niskanen, K. Harrabi, F. Yoshihara, Y. Nakamura, S. Lloyd, and J. S. Tsai. Quantum Coherent Tunable Coupling of Superconducting Qubits. *Science*, 316(5825):723–726, 2007. [1](#)
- [3] 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(7393):195–200, 2012. [1](#)
- [4] C. J. Ballance, V. M. Schäfer, J. P. Home, D. J. Szwer, S. C. Webster, D. T. C. Allcock, N. M. Linke, T. P. Harty, D. P. L. Aude Craik, D. N. Stacey, A. M. Steane, and D. M. Lucas. Hybrid quantum logic and a test of Bells inequality using two different atomic isotopes. *Nature*, 528(7582):384–386, 2015. [1](#)
- [5] D. Dimitropoulos and B. Jalali. Noise and information capacity in silicon nanophotonics. *IEEE Photonics Journal*, 7(3), 2015. [1](#), [118](#)
- [6] J. L. O’Brien, A. Furusawa, and J. Vučković. Photonic quantum technologies. *Nature Photonics*, 3(12):687–695, December 2009. [1](#), [7](#)
- [7] Y. Li, P. C. Humphreys, G. J. Mendoza, and S. C. Benjamin. Resource costs for fault-tolerant linear optical quantum computing. *Physical Review X*, 5(4):1–15, 2015. [1](#), [15](#), [18](#), [118](#)
- [8] J. P. Dowling, J. D. Franson, H. Lee, and G. J. Milburn. Towards scalable linear-optical quantum computers. *Quantum Information Processing*, 3(October):205–213, 2004. [1](#)
- [9] J. J. L. Morton, A. M. Tyryshkin, R. M. Brown, S. Shankar, B. W. Lovett, A. Ardavan, T. Schenkel, E. E. Haller, J. W. Ager, and S. A. Lyon. Solid-state quantum memory using the ^{31}P nuclear spin. *Nature*, 455(7216):1085–1088, 2008. [2](#)
- [10] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, M. J. Hatridge, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf. Quantum memory with millisecond coherence in circuit QED. *Physical Review B*, 94(1):1–8, 2016. [2](#)
- [11] F. Kaneda, F. Xu, J. Chapman, and P. G. Kwiat. Quantum-memory-assisted multi-photon generation for efficient quantum information processing. *Optica*, 4(9):1034–1037, 2017. [2](#), [8](#), [14](#), [15](#), [18](#), [72](#), [94](#), [121](#)
- [12] K. F. Reim, J. Nunn, X.-M. Jin, P. S. Michelberger, T. F. M. Champion, D. G. England, K. C. Lee, W. S. Kolthammer, N. K. Langford, and I. A. Walmsley. Multipulse Addressing of a Raman Quantum Memory: Configurable Beam Splitting and Efficient Readout. *Physical Review Letters*, 108(26):263602, June 2012. [2](#), [3](#), [15](#), [23](#)

- [13] G. T. Campbell, M. Hosseini, B. M. Sparkes, P. K. Lam, and B. C. Buchler. Time- and frequency-domain polariton interference. *New Journal of Physics*, 14, 2012. [2](#)
- [14] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982. [2](#), [9](#)
- [15] M. R. Sprague, P. S. Michelberger, T. F. M. Champion, D. G. England, J. Nunn, X.-M. Jin, W. S. Kolthammer, A. Abdolvand, P. St. J. Russell, and I. A. Walmsley. Broadband single-photon-level memory in a hollow-core photonic crystal fibre. *Nature Photonics*, 8(4):287–291, March 2014. [2](#), [5](#), [7](#), [24](#), [95](#), [96](#), [97](#), [120](#), [127](#)
- [16] 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 room-temperature Raman quantum memory. *New Journal of Physics*, 17(4):43006, 2015. [2](#), [5](#), [19](#), [24](#), [25](#), [55](#), [69](#), [72](#), [73](#), [80](#), [89](#), [119](#), [121](#)
- [17] 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(April):218–221, 2010. [4](#), [23](#), [30](#), [45](#), [58](#), [121](#), [123](#), [125](#)
- [18] R. Finkelstein, E. Poem, O. Michel, O. Lahad, and O. Firstenberg. Fast, noise-free memory for photon synchronization at room temperature, arXiv: 1708.01919 [quant-ph]. 2017. [5](#), [18](#), [26](#), [67](#), [69](#), [123](#), [124](#), [157](#)
- [19] M. Gündoğan, P. M. Ledingham, K. Kutluer, M. Mazzer, and H. de Riedmatten. Solid State Spin-Wave Quantum Memory for Time-Bin Qubits. *Physical Review Letters*, 114(23):1–5, 2015. [5](#), [71](#)
- [20] R. J. Glauber. The quantum theory of optical coherence. *Physical Review*, 130(6):2529–2539, 1963. [6](#), [76](#)
- [21] N. J. Cerf, O. Krüger, P. Navez, R. F. Werner, and M. M. Wolf. Non-Gaussian cloning of quantum coherent states is optimal. *Physical Review Letters*, 95(7), 2005. [6](#)
- [22] P. W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4):2493–2496, 1995. [6](#)
- [23] Joshua Nunn. *Quantum Memory in Atomic Ensembles*. PhD thesis, University of Oxford, 2008. [7](#), [28](#), [31](#), [32](#), [34](#), [35](#), [36](#), [37](#), [38](#), [39](#), [40](#), [41](#), [44](#), [51](#), [123](#), [156](#)
- [24] Z. Zheng, O. Mishina, N. Treps, and C. Fabre. Atomic quantum memory for multimode frequency combs. *Physical Review A - Atomic, Molecular, and Optical Physics*, 91(3):1–5, 2015. [7](#)
- [25] D. Stucki, N. Gisin, O. Guinnard, G. Ribordy, and H. Zbinden. Quantum Key Distribution over 67 km with a plug & play system. *Quantum Physics*, 4:8, 2002. [7](#), [9](#)
- [26] T. Zhong, J. M. Kindem, J. G. Bartholomew, J. Rochman, I. Craiciu, E. Miyazono, M. Bettinelli, E. Cavalli, V. Verma, S. W. Nam, F. Marsili, M. D. Shaw, A. D. Beyer, and A. Faraon. Nanophotonic rare-earth quantum memory with optically controlled retrieval. *Science*, 5959(August):eaan5959, 2017. [7](#), [22](#), [128](#)

- [27] P. Kumar, X. Li, P. L. Voss, J. Chen, and J. E. Sharping. Storage and long-distance distribution of telecommunications-band polarization entanglement generated in an optical fiber. *Optics Letters*, 30(10):1201–3, 2005. 8, 18
- [28] C. Xiong, X. Zhang, Z. Liu, M. J. Collins, A. Mahendra, L. G. Helt, M. J. Steel, D. Y. Choi, C. J. Chae, P. H. W. Leong, and B. J. Eggleton. Active temporal multiplexing of indistinguishable heralded single photons. *Nature Communications*, 7:10853, 2016. 8
- [29] D. V. Vasilyev, I. V. Sokolov, and E. S. Polzik. Quantum memory for images: A quantum hologram. *Physical Review A - Atomic, Molecular, and Optical Physics*, 77(2), 2008. 8
- [30] M. Fujiwara, A. Waseda, R. Nojima, S. Moriai, W. Ogata, and M. Sasaki. Unbreakable distributed storage with quantum key distribution network and password-authenticated secret sharing. *Scientific Reports*, 6(1):28988, 2016. 8
- [31] 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(7533):177–180, 2015. 8, 20
- [32] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden. Quantum cryptography. *Reviews of Modern Physics*, 74(1):145–195, 2002. 9
- [33] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984. 9
- [34] A. K. Ekert. Quantum cryptography based on Bell’s theorem. *Physical Review Letters*, 67(6):661–663, 1991. 9
- [35] J. Yin, Y. Cao, Y.-H. Li, S.-K. Liao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, B. Li, H. Dai, G.-B. Li, Q.-M. Lu, Y.-H. Gong, Y. Xu, S.-L. Li, F.-Z. Li, Y.-Y. Yin, Z.-Q. Jiang, M. Li, J.-J. Jia, G. Ren, D. He, Y.-L. Zhou, X.-X. Zhang, N. Wang, X. Chang, Z.-C. Zhu, N.-L. Liu, Y.-A. Chen, C.-Y. Lu, R. Shu, C.-Z. Peng, J.-Y. Wang, and J.-W. Pan. Satellite-based entanglement distribution over 1200 kilometers. *Science*, 356(6343):1140–1144, 2017. 9
- [36] W. A. Gambling. The rise and rise of optical fibers. *IEEE Journal on Selected Topics in Quantum Electronics*, 6(6):1084–1093, 2000. 9, 99
- [37] 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(26):5932–5935, 1998. 9
- [38] L. M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller. Long-distance quantum communication with atomic ensembles and linear optics. *Nature*, 414(6862):413–8, November 2001. 11, 118
- [39] C. V. Raman and K. S. Krishnan. A New Type of Secondary Radiation. *Nature*, 121(3048):501–502, 1928. 11
- [40] C. W. Chou, H. de Riedmatten, D. Felinto, S. V. Polyakov, S. J. Van Enk, and H. J. Kimble. Measurement induced entanglement for excitation stored in remote atomic ensembles. *Nature*, 438(December):828–832, 2006. 11

- [41] Y.-W. Cho, G. T. Campbell, J. L. Everett, J. Bernu, D. B. Higginbottom, M. T. Cao, J. Geng, N. P. Robins, P. K. Lam, and B. C. Buchler. Highly efficient optical quantum memory with long coherence time in cold atoms. *Optica*, 3(1):100, 2016. [11](#), [21](#), [123](#)
- [42] C. Laplane, P. Jobez, J. Etesse, N. Gisin, and M. Afzelius. Multimode and Long-Lived Quantum Correlations between Photons and Spins in a Crystal. *Physical Review Letters*, 118(21):1–5, 2017. [11](#), [22](#), [24](#), [92](#)
- [43] D. E. Chang, V. Vuletić, and M. D. Lukin. Quantum nonlinear optics photon by photon. *Nature Photonics*, 8(September):685–694, August 2014. [11](#)
- [44] T. G. Tiecke, J. D. Thompson, N. P. de Leon, L. R. Liu, V. Vuletić, and M. D. Lukin. Nanophotonic quantum phase switch with a single atom. *Nature*, 508(7495):241–244, 2014. [11](#), [16](#)
- [45] A. Sipahigil, R. E. Evans, D. D. Sukachev, M. J. Burek, J. Borregaard, M. K. Bhaskar, C. T. Nguyen, J. L. Pacheco, H. a. Atikian, C. Meuwly, R. M. Camacho, F. Jelezko, E. Bielejec, H. Park, M. Lončar, and M. D. Lukin. An integrated diamond nanophotonics platform for quantum-optical networks. *Science*, 354(6314):847–850, 2016. [11](#), [19](#)
- [46] T. Peyronel, O. Firstenberg, Q.-Y. Liang, S. Hofferberth, A. V Gorshkov, T. Pohl, M. D Lukin, and V. Vuletić. Quantum nonlinear optics with single photons enabled by strongly interacting atoms. *Nature*, 488(7409):57–60, August 2012. [11](#)
- [47] P. Kok, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn. Linear optical quantum computing with photonic qubits. *Reviews of Modern Physics*, 79(1):135–174, January 2007. [11](#)
- [48] M. A. Nielsen. Cluster-state quantum computation. *Reports on Mathematical Physics*, 57(1):147–161, 2006. [12](#)
- [49] E. Knill, R. Laflamme, and G. J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409(6816):46–52, January 2001. [11](#)
- [50] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger. Experimental quantum teleportation. *Nature*, 390(6660):575–579, 1997. [12](#)
- [51] H. Ishio, J. Minowa, and K. Nosu. Review and status of wavelength-division-multiplexing technology and its application. *Journal of Lightwave Technology*, 2(4):3–9, 1984. [12](#)
- [52] C. Joshi, A. Farsi, S. Clemmen, S. Ramelow, and A. L. Gaeta. Frequency Multiplexing for Quasi-Deterministic Heralded Single-Photon Sources, arXiv:1707.00048 [quant-ph]. 2017. [12](#)
- [53] R. J. A. Francis-Jones, R. A. Hoggarth, and P. J. Mosley. All-fiber multiplexed source of high-purity single photons. *Optica*, 3(11):1270–1273, 2016. [12](#)
- [54] T. B. Pittman, B. C. Jacobs, and J. D. Franson. Single photons on pseudodemand from stored parametric down-conversion. *Physical Review A*, 66(4):042303, 2002. [13](#)

- [55] 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(13):133601, March 2013. [13](#), [15](#), [119](#), [123](#)
- [56] A. Christ and C. Silberhorn. Limits on the deterministic creation of pure single-photon states using parametric down-conversion. *Physical Review A - Atomic, Molecular, and Optical Physics*, 85(2):1–6, 2012. [13](#), [73](#)
- [57] N. Somaschi, V. Giesz, L. De Santis, J. C. Loredó, M. P. Almeida, G. Hornecker, S. L. Portalupi, T. Grange, C. Antón, J. Demory, C. Gómez, I. Sagnes, N. D. Lanzillotti-Kimura, a. Lemaître, A. Auffeves, A. G. White, L. Lanco, and P. Senellart. Near-optimal single-photon sources in the solid state. *Nature Photonics*, 10(5):340–345, 2016. [13](#), [14](#)
- [58] M. Dabrowski, M. Parniak, and W. Wasilewski. Einstein-Podolsky-Rosen paradox in a hybrid bipartite system. *Optica*, 4(2):272–275, 2017. [13](#), [17](#)
- [59] M. Parniak, M. Dabrowski, M. Mazelanik, A. Leszczyński, M. Lipka, and W. Wasilewski. Wavevector multiplexed quantum memory as a universal platform for quantum state preparation, arXiv: 1706.04426 [quant-ph]. 2017. [13](#)
- [60] 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(6):361–365, 2017. [14](#)
- [61] D. G. England, K. A. G. Fisher, J.-P. W. MacLean, P. J. Bustard, K. Heshami, K. J. Resch, and B. J. Sussman. Phonon-Mediated Nonclassical Interference in Diamond. *Physical Review Letters*, 117(7):1–5, 2016. [15](#), [17](#), [23](#), [126](#)
- [62] B. Brecht, A. Eckstein, R. Ricken, V. Quiring, H. Suche, L. Sansoni, and C. Silberhorn. Demonstration of coherent time-frequency Schmidt mode selection using dispersion-engineered frequency conversion. *Physical Review A - Atomic, Molecular, and Optical Physics*, 90(3):1–5, 2014. [15](#), [16](#), [23](#), [122](#)
- [63] M. Hosseini, B. M. Sparkes, G. Hétet, J. J. Longdell, P. K. Lam, and B. C. Buchler. Coherent optical pulse sequencer for quantum applications. *Nature*, 461(7261):241–245, 2009. [15](#), [21](#)
- [64] P. C. Humphreys, W. S. Kolthammer, J. Nunn, M. Barbieri, A. Datta, and I. A. Walmsley. Continuous-Variable Quantum Computing in Optical Time-Frequency Modes using Quantum Memories. *Physical Review Letters*, 130502(September):5, 2014. [16](#)
- [65] E. Saglamyurek, N. Sinclair, J. A. Slater, K. Heshami, D. Oblak, and W. Tittel. An integrated processor for photonic quantum states using a broadband light-matter interface. *New Journal of Physics*, 16(6):065019, June 2014. [16](#), [22](#)
- [66] 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*, 7:11200, 2016. [16](#), [23](#), [89](#)

- [67] R. Löw, H. Weimer, J. Nipper, J. B. Balewski, B. Butscher, H. P. Büchler, and T. Pfau. An experimental and theoretical guide to strongly interacting Rydberg gases. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(11):113001, June 2012. 16, 25
- [68] M. Kiffner, A. Feizpour, K. T. Kaczmarek, D. Jaksch, and J. Nunn. Two-way interconversion of millimeter-wave and optical fields in Rydberg gases. *New Journal of Physics*, 18(9), 2016. 16
- [69] J. Han, T. Vogt, C. Gross, D. Jaksch, M. Kiffner, and W. Li. Free-space microwave-to-optical conversion via six-wave mixing in Rydberg atoms, arXiv: 1701.07969 [quant-ph]. pages 1–6, 2017. 16
- [70] D. Maxwell, D. Szwer, D. Paredes-Barato, H. Busche, J. Pritchard, a. Gauguet, K. Weatherill, M. Jones, and C. S. Adams. Storage and Control of Optical Photons Using Rydberg Polaritons. *Physical Review Letters*, 110(10):103001, March 2013. 16, 25
- [71] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. D. Yavuz, T. G. Walker, and M. Saffman. Observation of Rydberg blockade between two atoms. *Nature Physics*, 5(2):110–114, January 2009. 16
- [72] V. Parigi, E. Bimbard, J. Stanojevic, A. J. Hilliard, F. Nogrette, R. Tualle-Brouri, A.i Ourjoumtsev, and P. Grangier. Observation and Measurement of Interaction-Induced Dispersive Optical Nonlinearities in an Ensemble of Cold Rydberg Atoms. *Physical Review Letters*, 109(23), 2012. 16
- [73] D. Tiarks, S. Schmidt, G. Rempe, and S. Dürr. Optical π phase shift created with a single-photon pulse. *Science Advances*, 2(4):e1600036–e1600036, 2016. 17
- [74] J. D. Thompson, T. L. Nicholson, Q.-Y. Liang, S. H. Cantu, A. V. Venkatramani, S. Choi, I. A. Fedorov, D. Viscor, T. Pohl, M. D. Lukin, and V. Vuletić. Symmetry-protected collisions between strongly interacting photons. *Nature*, 542(7640):206–209, 2017. 17
- [75] D. Jaksch, J. I. Cirac, P. Zoller, and S. L. Rolston. Fast quantum gates for neutral atoms. *Physical Review Letters*, 5(1):5–8, 2000. 17
- [76] O. Firstenberg, T. Peyronel, Q.-Y. Liang, A. V. Gorshkov, M. D. Lukin, and V. Vuletić. Attractive photons in a quantum nonlinear medium. *Nature*, 502(7469):71–5, October 2013. 17
- [77] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenbergh, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson. Experimental loophole-free violation of a Bell inequality using entangled electron spins separated by 1.3 km. *Nature*, 526, 2015. 17, 18
- [78] K. C. Lee, B. J. Sussman, M. R. Sprague, P. Michelberger, K. F. Reim, J. Nunn, N. K. Langford, P. J. Bustard, D. Jaksch, and I. A. Walmsley. Macroscopic non-classical states and terahertz quantum processing in room-temperature diamond. *Nature Photonics*, 6(1):41–44, 2011. 17

- [79] K. C. Lee, M. R. Sprague, B. J. Sussman, J. Nunn, N. K. Langford, X.-M. Jin, T. Champion, P. Michelberger, K. F. Reim, D. England, D. Jaksch, and I. A. Walmsley. Entangling Macroscopic Diamonds at Room Temperature. *Science*, 334(6060):1253–1256, 2011. 17
- [80] A. Tiranov, S. Designolle, Emmanuel Z. Cruzeiro, J. Lavoie, N. Brunner, M. Afzelius, M. Huber, and N. Gisin. Quantification of multi-dimensional photonic entanglement stored in a quantum memory based on sparse data. pages 1–9, 2016. 17
- [81] A. Tiranov, J. Lavoie, P. C. Strassmann, N. Sangouard, M. Afzelius, F. Bussi eres, and N. Gisin. Demonstration of Light-Matter Micro-Macro Quantum Correlations. *Physical Review Letters*, 116(19):1–6, 2016. 17
- [82] P. Zarkeshian, C. Deshmukh, N. Sinclair, S. K. Goyal, G. H. Aguilar, P. Lefebvre, M. Grimaud Puigibert, V. B. Verma, F. Marsili, M. D. Shaw, S. W. Nam, K. Heshami, D. Oblak, W. Tittel, and C. Simon. Entanglement between more than two hundred macroscopic atomic ensembles in a solid, arXiv:1703.04709 [quant-ph]. pages 1–10, 2017. 17, 22
- [83] F. Fr owis, P. C. Strassmann, A. Tiranov, C. Gut, J. Lavoie, N. Brunner, F. Bussi eres, M. Afzelius, and N. Gisin. Experimental certification of millions of genuinely entangled atoms in a solid, arXiv:1703.04704 [quant-ph]. (1):1–11, 2017. 17
- [84] S. L. Christensen, J. B. B egu in, E. Bookjans, H. L. S oren sen, J. H. M uller, J. Appel, and E. S. Polzik. Quantum interference of a single spin excitation with a macroscopic atomic ensemble. *Physical Review A - Atomic, Molecular, and Optical Physics*, 89(3):1–7, 2014. 17, 126
- [85] C. Clausen, I. Usmani, F. Bussi eres, N. Sangouard, M. Afzelius, H. de Riedmatten, and N. Gisin. Quantum storage of photonic entanglement in a crystal. *Nature*, 469(7331):508–11, January 2011. 18, 125, 126
- [86] D. G. England, P. S. Michelberger, T. F. M. Champion, K. F. Reim, K. C. Lee, M. R. Sprague, X.-M. Jin, N. K. Langford, W. S. Kolthammer, J. Nunn, and I. A. Walmsley. High-fidelity polarization storage in a gigahertz bandwidth quantum memory. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(12):124008, June 2012. 18, 24, 70, 157
- [87] B. Julsgaard, J. Sherson, J. I. Cirac, J. Fiur asek, and E. S. Polzik. Experimental demonstration of quantum memory for light. *Nature*, 432(7016):482–486, 2004. 18
- [88] V. Damon, M. Bonarota, A. Louchet-Chauvet, T. Chaneli ere, and J. L. Le Gou et. Revival of silenced echo and quantum memory for light. *New Journal of Physics*, 13, 2011. 18
- [89] H. de Riedmatten, M. Afzelius, M. U. Staudt, C. Simon, and N. Gisin. A solid-state lightmatter interface at the single-photon level. *Nature*, 456(7223):773–777, 2008. 18, 21
- [90] K. Makino, Y. Hashimoto, J.-I. Yoshikawa, H. Ohdan, T. Toyama, P. van Loock, and A. Furusawa. Synchronization of optical photons for quantum information processing. *Science Advances*, (May):1–7, 2015. 18

- [91] H. P. Specht, C. Nölleke, A. Reiserer, M. Uphoff, E. Figueroa, S. Ritter, and G. Rempe. A single-atom quantum memory. *Nature*, 473(7346):190–193, 2011. 19
- [92] G. D. Fuchs, G. Burkard, P. V. Klimov, and D. D. Awschalom. A quantum memory intrinsic to single nitrogenvacancy centres in diamond. *Nature Physics*, 7(10):789–793, 2011. 19
- [93] D. D. Sukachev, A. Sipahigil, C. T. Nguyen, M. K. Bhaskar, R. E. Evans, F. Jelezko, and M. D. Lukin. The silicon-vacancy spin qubit in diamond: quantum memory exceeding ten milliseconds and single-shot state readout, arXiv: 1708.08852 [quant-ph]. 2017. 19
- [94] S. L. Mouradian, T. Schröder, C. B. Poitras, L. Li, J. Goldstein, E. H. Chen, M. Walsh, J. Cardenas, M. L. Markham, D. J. Twitchen, M. Lipson, and D. Englund. Scalable integration of long-lived quantum memories into a photonic circuit. *Physical Review X*, 5(3), 2015. 19
- [95] K. T. Kaczmarek, D. J. Saunders, M. R. Sprague, W. S. Kolthammer, A. Feizpour, P. M. Ledingham, B. Brecht, E. Poem, I. A. Walmsley, and J. Nunn. Ultrahigh and persistent optical depths of caesium in Kagomé-type hollow-core photonic crystal fibres. *Optics Letters*, 40(23):005582, 2015. 19, 27, 102, 120, 128
- [96] S. 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(17):173601, October 2001. 20
- [97] 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(2):1–4, 2008. 20, 21
- [98] M. Hosseini, B. M. Sparkes, G. T. Campbell, P. K. Lam, and B. C. Buchler. Storage and manipulation of light using a Raman gradient-echo process. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(12):124004, 2012. 20, 122
- [99] G. Hétet, M. Hosseini, B. M. Sparkes, D. Oblak, B. C. Buchler, and P. K. Lam. Photon echoes generated by reversing magnetic field gradients in a rubidium vapor. *Optics letters*, 33(20):2323–5, 2008. 20
- [100] 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*, 063034, 2017. 20, 25, 41, 48, 53, 55, 57, 58, 99, 123
- [101] A. L. Alexander, J. J. Longdell, M. J. Sellars, and N. B. Manson. Photon Echoes Produced by Switching Electric Fields. *Physical review letters*, 94(1):94–97, 2006. 20
- [102] 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, arXiv:1611.04315 [quant-ph]. 2016. 20
- [103] D. A Steck. Cesium D Line Data Cesium Physical and Optical Properties, accessible at: <http://steck.us/alkalidata/>, 2010. 21, 102, 149

- [104] Guokui Liu and Bernard Jacquier. *Spectroscopic properties of rare earths in optical materials*. Springer, 2010. [21](#)
- [105] M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler. Unconditional room-temperature quantum memory. *Nature Physics*, 7(10):794–798, 2011. [21](#), [121](#)
- [106] G. Hétet, J. Longdell, M. Sellars, P. Lam, and B. Buchler. Multimodal Properties and Dynamics of Gradient Echo Quantum Memory. *Physical Review Letters*, 101(20):203601, November 2008. [21](#), [122](#)
- [107] 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, 2013. [21](#)
- [108] M. Hosseini, B.M. Sparkes, G. Campbell, P.K. Lam, and B.C. Buchler. High efficiency coherent optical memory with warm rubidium vapour. *Nature Communications*, 2(2):174, February 2011. [21](#)
- [109] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin. Multimode quantum memory based on atomic frequency combs. *Physical Review A*, 79(5):052329, May 2009. [21](#)
- [110] M. Afzelius, I. Usmani, A. Amari, B. Lauritzen, A. Walther, C. Simon, N. Sangouard, J. Miná, H. de Riedmatten, N. Gisin, and S. Kröll. Demonstration of atomic frequency comb memory for light with spin-wave storage. *Physical Review Letters*, 104(4):1–4, 2010. [21](#), [22](#), [30](#)
- [111] M. Sabooni, Q. Li, S. Kröll, and L. Rippe. Efficient quantum memory using a weakly absorbing sample. *Physical Review Letters*, 110(13):1–5, 2013. [22](#)
- [112] A. Tiranov, J. Lavoie, A. Ferrier, P. Goldner, V. B. Verma, S. W. Nam, R. P. Mirin, A. E. Lita, F. Marsili, H. Herrmann, C. Silberhorn, N. Gisin, M. Afzelius, and F. Bussières. Storage of hyperentanglement in a solid-state quantum memory. *Optica*, 2(4):279, 2015. [22](#)
- [113] 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, 2014. [22](#)
- [114] A. Seri, A. Lenhard, D. Rieländer, M. Gündoğan, P. M. Ledingham, M. Mazzera, and H. de Riedmatten. Quantum correlations between single telecom photons and a multimode on-demand solid state quantum memory. *Physical Review X*, 7(021028):1–13, 2017. [22](#), [77](#), [92](#)
- [115] K. Kutluer, M. Mazzera, and H. de Riedmatten. Solid-State Source of Nonclassical Photon Pairs with Embedded Multimode Quantum Memory. *Physical Review Letters*, 118(21):1–6, 2017. [22](#), [122](#)
- [116] M. Gündoğan, M. Mazzera, P. M. Ledingham, M. Cristiani, and H. de Riedmatten. Coherent storage of temporally multimode light using a spin-wave atomic frequency comb memory. *New Journal of Physics*, 15(4):045012, April 2013. [22](#), [123](#)
- [117] A. E. Kozhekin, K. Molmer, and E. Polzik. Quantum Memory for Light. *Physical Review A*, 62(February):1–5, 2000. [22](#)

- [118] M. Fleischhauer and M. D. Lukin. Dark-state polaritons in electromagnetically induced transparency. *Physical review letters*, 84(22):5094–7, May 2000. 22
- [119] M. Fleischhauer and M. D. Lukin. Quantum memory for photons: Dark-state polaritons. *Physical Review A. Atomic, Molecular, and Optical Physics*, 65(2):1–12, 2002. 22
- [120] C. Liu, Z. Dutton, C. H. Behroozi, and L. V. Hau. Observation of coherent optical information storage in an atomic medium using halted light pulses. *Nature*, 409(6819):490–493, 2001. 22, 23
- [121] K.-J. Boller, A. Imamoglu, and S. E. Harris. Observation of electromagnetically induced transparency. *Physical Review Letters*, 66(20):2593–2596, 1991. 22
- [122] D. F. Phillips, A. Fleischhauer, A. Mair, R. L. Walsworth, and M. D. Lukin. Storage of Light in Atomic Vapor. *Physical Review Letters*, 86(5):783–786, January 2001. 23
- [123] A. V. Turukhin, V. S. Sudarshanam, M. S. Shahriar, J. A. Musser, B. S. Ham, and P. R. Hemmer. Observation of ultraslow and stored light pulses in a solid. *Physical Review Letters*, 88(2):1–4, 2002. 23
- [124] Y. O. Dudin, L. Li, and A. Kuzmich. Light storage on the time scale of a minute. *Physical Review A - Atomic, Molecular, and Optical Physics*, 87(3):1–4, 2013. 23
- [125] 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(3):033601, July 2013. 23, 121
- [126] 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(8):1–5, 2013. 23, 121
- [127] D. Schraft, M. Hain, N. Lorenz, and T. Halfmann. Stopped Light at High Storage Efficiency in a Pr^{3+} : Y_2SiO_5 Crystal. *Physical Review Letters*, 116(7):1–6, 2016. 23
- [128] M. D. Eisaman, L. Childress, A. Andre, F. Massou, A. S. Zibrov, and M. D. Lukin. Shaping quantum pulses of light via coherent atomic memory. *Physical Review Letters*, 93(23):1–4, 2004. 23
- [129] T. Chanelière, D.N. Matsukevich, S.D. Jenkins, S.Y. Lan, T.A.B. Kennedy, and A. Kuzmich. Storage and retrieval of single photons transmitted between remote quantum memories. *Nature*, 438(7069):833, 2005. 23
- [130] M. D. Eisaman, A. André, F. Massou, M. Fleischhauer, A. S. Zibrov, and M. D. Lukin. Electromagnetically induced transparency with tunable single-photon pulses. *Nature*, 438(7069):837–41, 2005. 23
- [131] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen. Photon storage in Λ -type optically dense atomic media. III. Effects of inhomogeneous broadening. *Physical Review A*, 76(3):033806, 2007. 23
- [132] P. J. Bustard, R. Lausten, D. G. England, and B. J. Sussman. Toward quantum processing in molecules: A THz-bandwidth coherent memory for light. *Physical Review Letters*, 111(8):1–5, 2013. 23

- [133] D. G. England, K. A.G. 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(5):053602, 2015. [23](#), [24](#), [121](#)
- [134] J. Nunn, I. A. Walmsley, M. Raymer, K. Surmacz, F. Waldermann, Z. Wang, and D. Jaksch. Mapping broadband single-photon wave packets into an atomic memory. *Physical Review A*, 75(1):011401, January 2007. [23](#), [41](#), [44](#), [120](#)
- [135] N. V. Vitanov, A. A. Rangelov, B. W. Shore, and K. Bergmann. Stimulated Raman adiabatic passage in physics, chemistry, and beyond. *Reviews of Modern Physics*, 89(1):1–66, 2017. [23](#)
- [136] 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(26):1–4, 2008. [23](#), [41](#), [122](#)
- [137] G. T. Campbell, O. Pinel, M. Hosseini, T. C. Ralph, B. C. Buchler, and P. K. Lam. Configurable unitary transformations and linear logic gates using quantum memories. *Physical Review Letters*, 113(6):1–5, 2014. [23](#)
- [138] D.-S. Ding, W. Zhang, Z.-Y. Zhou, S. Shi, B.-S. Shi, and G.-C. Guo. Raman quantum memory of photonic polarized entanglement. *Nature Photonics*, 9(5):332–338, 2015. [24](#), [26](#)
- [139] 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, Dileep V. Reddy, M. G. Raymer, and I. A. Walmsley. Theory of noise suppression in Λ -type quantum memories by means of a cavity. *Physical Review A*, 96(1):1–12, 2017. [24](#), [25](#), [34](#), [70](#)
- [140] 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. A Cavity-Enhanced Room-Temperature Broadband Raman Memory. *Physical Review Letters*, 3(March):25–27, 2015. [24](#), [25](#), [41](#), [45](#), [47](#), [49](#), [70](#), [120](#), [122](#), [123](#)
- [141] G. Romanov, C. O'Brien, and I. Novikova. Suppression of the four-wave mixing amplification via Raman absorption. *Journal of Modern Optics*, 00(00):1–10, 2016. [25](#), [41](#)
- [142] 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(5):053603, July 2011. [25](#)
- [143] E. Distante, P. Farrera, A. Padrón-Brito, D. Paredes-Barato, G. Heinze, and H. de Riedmatten. Storing single photons emitted by a quantum memory on a highly excited Rydberg state. *Nature Communications*, 8:14072, 2017. [25](#)
- [144] N. Šibalić, J. M. Kondo, C. S. Adams, and K. J. Weatherill. Dressed-state electromagnetically induced transparency for light storage in uniform phase spin-waves. pages 1–7, 2016. [25](#)
- [145] B. Zhao, Y.-A. Chen, X.-H. Bao, T. Strassel, C.-S. Chu, 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(2):95–99, 2009. [25](#), [65](#), [124](#)

- [146] 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. A room-temperature noise-free quantum memory for broadband light, arXiv: 1704.00013 [quant-ph]. pages 1–12, 2017. 26, 27, 28, 45, 55, 120, 121
- [147] J. E. Sansonetti. Wavelengths, transition probabilities, and energy levels for the spectra of rubidium (Rb i through Rb XXXVII). *Journal of Physical and Chemical Reference Data*, 35(1):301–421, 2006. 26
- [148] M. G. Raymer and J. L. Carlsten. Simultaneous observations of stimulated raman scattering and stimulated collision-induced fluorescence. *Physical Review Letters*, 39(21):1326–1329, 1977. 26
- [149] Theresa F. M. Champion. *Towards storage and retrieval of non-classical light in a broadband quantum memory - An investigation of free-space and cavity Raman memories*. PhD thesis, University of Oxford, 2015. 28, 58, 148, 155
- [150] G. S. Thekkadath, K. Heshami, D. G. England, P. Bustard, B. J. Sussman, and M. Spanner. Optical quantum memory for ultrafast photons using molecular alignment. *Journal of Modern Optics*, 0340(i):1–6, 2015. 29
- [151] W. Grice and I. A. Walmsley. Spectral information and distinguishability in type-II down-conversion with a broadband pump. *Phys. Rev. A*, 56(2):1627–1634, 1997. 31, 74
- [152] R. Loudon. *The Quantum Theory of Light*. OUP, Oxford, 2000. 32
- [153] J. Kołodyński, J. Chwedeńczuk, and W. Wasilewski. Eigenmode description of Raman scattering in atomic vapors in the presence of decoherence. *Physical Review A - Atomic, Molecular, and Optical Physics*, 86(1), 2012. 37
- [154] 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(12):1–4, 2007. 41, 51
- [155] E. Suli and D. Mayers. *An introduction to numerical analysis*. 2003. 45
- [156] O. Lazo-Arjona. A repository for solvers for the Maxwell-Bloch equations of various quantum memories, and calculations using them, accessible at: https://github.com/oscarlazoarjona/quantum_memories, 2017. 47
- [157] L. N. Trefethen. *Spectral Methods in MATLAB*. SIAM, Philadelphia, 2000. 48
- [158] J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, 1965. 53
- [159] L. Weller, R. J. Bettles, P. Siddons, C. S. Adams, and I. G. Hughes. Absolute absorption on the rubidium D1 line including resonant dipole-dipole interactions. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 195006(44):0–5, 2011. 53, 150
- [160] L. Essen and J. V. L. Parry. An atomic standard of frequency and time interval: a caesium resonator. *Nature*, 1955. 57

- [161] B. Huber, T. Baluktsian, M. Schlagmüller, A. Kölle, H. Kübler, R. Löw, and T. Pfau. GHz Rabi Flopping to Rydberg States in Hot Atomic Vapor Cells. *Physical Review Letters*, 107(24):243001, December 2011. 57
- [162] J. E. Sansonetti. Wavelengths, transition probabilities, and energy levels for the spectra of cesium (Cs I-Cs LV). *Journal of Physical and Chemical Reference Data*, 38(4):761–923, 2009. 57
- [163] M. S. Safronova, U. I. Safronova, and Charles W. Clark. Magic wavelengths, matrix elements, polarizabilities, and lifetimes of Cs. *Physical Review A*, 94(1):1–10, 2016. 57
- [164] G. A. Pitz, C. D. Fox, and G. P. Perram. Pressure broadening and shift of the cesium D2 transition by the noble gases and N₂, H₂, HD, D₂, CH₄, C₂H₆, CF₄, and He₃ with comparison to the D1 transition. *Physical Review A - Atomic, Molecular, and Optical Physics*, 82(4):1–9, 2010. 59
- [165] R. Hanbury Brown and R. Q. Twiss. Correlation between photons in two coherent beams of light. *Nature*, 15(177):27–29, 1956. 62
- [166] N. Ph. Georgiades, E. S. Polzik, and H. J. Kimble. Two-photon spectroscopy of the 6S 1/2 -6D transition of trapped atomic cesium. *Optics Letters*, 19(18):1474–1476, 1994. 65
- [167] A. Urvoy, C. Carr, R. Ritter, C. S. Adams, K. J. Weatherill, and R. Löw. Optical coherences and wavelength mismatch in ladder systems. 2013. 65
- [168] O. Lazo-Arjona. FAST Atomic Spectroscopy from Theory, accessible at: <https://github.com/oscarlazoarjona/fast>, 2017. 69
- [169] S. Manz, T. Fernholz, J. Schmiedmayer, and J.-W. Pan. Collisional decoherence during writing and reading quantum states. *Physical Review A - Atomic, Molecular, and Optical Physics*, 75(4):1–4, 2007. 70
- [170] P. S. Michelberger, M. Karpiński, I. A. Walmsley, and J. Nunn. Optimal interfacing a GHz-bandwidth heralded single photon source with on-demand, broadband quantum memories, arXiv: 1709.03155 [quant-ph]. 2017. 73
- [171] D. C. Burnham and D. L. Weinberg. Observation of simultaneity in parametric production of optical photon pairs. *Physical Review Letters*, 25(2):84–87, 1970. 73
- [172] M. Takeoka, R.-B. Jin, and M. Sasaki. Full analysis of multi-photon pair effects in spontaneous parametric down conversion based photonic quantum information processing. *New J. Phys.*, 17:0–17, 2014. 74
- [173] A. B. U'Ren, C. Silberhorn, K. Banaszek, I. A. Walmsley, R. Erdmann, W. P. Grice, and M. G. Raymer. Generation of Pure-State Single-Photon Wavepackets by Conditional Preparation Based on Spontaneous Parametric Downconversion. *Laser Phys.*, 15(146):1–14, 2006. 75
- [174] 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(7), 2015. 75
- [175] D N Klyshko. Use of two-photon light for absolute calibration of photoelectric detectors. *Soviet Journal of Quantum Electronics*, 10(9):1112–1117, 1980. 75

- [176] P. Sekatski, N. Sangouard, F. Bussi eres, C. Clausen, N. Gisin, and H. Zbinden. Detector imperfections in photon-pair source characterization. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(12):124016, 2012. 77, 90, 91, 94
- [177] P. J. Mosley, A. Christ, A. Eckstein, and C. Silberhorn. Direct measurement of the spatial-spectral structure of waveguided parametric down-conversion. *Physical Review Letters*, 103(23):1–4, 2009. 78
- [178] M. Fiorentino, S. M. Spillane, R. G. Beausoleil, T. D. Roberts, P. Battle, and M. W. Munro. Spontaneous parametric down-conversion in periodically poled KTP waveguides and bulk crystals. *Optics Express*, 15(12):7479, 2007. 79
- [179] H. Vanherzeele and J. D. Bierlein. Magnitude of the nonlinear-optical coefficients of KTiOPO4. *Optics letters*, 17(14):982–984, 1992. 79
- [180] A. Christ, K. Laiho, A. Eckstein, K. N. Cassemiro, and C. Silberhorn. Probing multimode squeezing with correlation functions. *New Journal of Physics*, 13, 2011. 91
- [181] F. Benabid, J. C. Knight, G. Antonopoulos, and P. St. J. Russell. Stimulated Raman scattering in hydrogen-filled hollow-core photonic crystal fiber. *Science (New York, N. Y.)*, 298(5592):399–402, 2002. 96
- [182] V. Venkataraman, K. Saha, and A. L. Gaeta. Phase modulation at the few-photon level for weak-nonlinearity-based quantum computing. *Nature Photonics*, 7(February 2013):5–8, 2013. 96
- [183] G. Epple, K. S. Kleinbach, T. G. Euser, N. Y. Joly, T. Pfau, P. St. J. Russell, and R. L  w. Rydberg atoms in hollow-core photonic crystal fibres. *Nature communications*, 5(May):4132, January 2014. 96
- [184] M. R. Sprague, D. G. England, A. Abdolvand, J. Nunn, X.-M. Jin, W. S. Kolthammer, M. Barbieri, B. Rigal, P. S. Michelberger, P. St. J. Champion, T. F. M. and Russell, and I. A. Walmsley. Efficient optical pumping and high optical depth in a hollow-core photonic-crystal fibre for a broadband quantum memory. *New Journal of Physics*, 15(5):055013, May 2013. 96, 102, 104, 106, 127
- [185] M. Meucci, E. Mariotti, P. Bicchi, C. Marinelli, and L. Moi. Light-Induced Atom Desorption. *Europhysics Letters (EPL)*, 25(9):639–643, March 1994. 96, 101
- [186] B. E. A. Saleh and M. C. Teich. *Fundamentals of Photonics*, 2nd Edition, 2007. 97
- [187] R. F. Cregan, B. J. Mangan, J. C. Knight, T. A. Birks, P. S. Russell, P. J. Roberts, and D. C. Allan. Single-Mode Photonic Band Gap Guidance of Light in Air. *Science*, 285(5433):1537–1539, 1999. 97
- [188] F. Couny, F. Benabid, and P. S. Light. Large-pitch kagome-structured hollow-core photonic crystal fiber. *Optics letters*, 31(24):3574–3576, 2006. 97
- [189] S.-J. Im, A. Husakou, and J. Herrmann. Guiding properties and dispersion control of kagome lattice hollow-core photonic crystal fibers. *Optics express*, 17(15):13050–13058, 2009. 97
- [190] C. Schori, B. Julsgaard, J. L. S  rensen, and E. S. Polzik. Recording Quantum Properties of Light in a Long-Lived Atomic Spin State: Towards Quantum Memory. *Physical Review Letters*, 89(5):057903, 2002. 97

- [191] S. J. Seltzer, D. J. Michalak, M. H. Donaldson, M. V. Balabas, S. K. Barber, S. L. Bernasek, M. A. Bouchiat, A. Hexemer, A. M. Hibberd, D. F. Jackson Kimball, C. Jaye, T. Karaulanov, F. A. Narducci, S. A. Rangwala, H. G. Robinson, a. K. Shmakov, D. L. Voronov, V. V. Yashchuk, A. Pines, and D. Budker. Investigation of antirelaxation coatings for alkali-metal vapor cells using surface science techniques. *Journal of Chemical Physics*, 133(14), 2010. [97](#)
- [192] A. Bhagwat, A. Slepko, V. Venkataraman, P. Londero, and A. Gaeta. On-demand all-optical generation of controlled Rb-vapor densities in photonic-band-gap fibers. *Physical Review A*, 79(6):063809, June 2009. [97](#), [106](#), [107](#), [108](#)
- [193] Michael R. Sprague. *Raman Memory for Entanglement in Diamonds and Light Storage in Optical Fibres*. PhD thesis, University of Oxford, 2014. [99](#), [153](#), [154](#)
- [194] T. Karaulanov, M. T. Graf, D. English, S. M. Rochester, Y. J. Rosen, K. Tsigutkin, D. Budker, E. B. Alexandrov, M. V. Balabas, D. F. Jackson Kimball, F. A. Narducci, S. Pustelny, and V. V. Yashchuk. Controlling atomic vapor density in paraffin-coated cells using light-induced atomic desorption. *Physical Review A - Atomic, Molecular, and Optical Physics*, 79(1):1–9, 2009. [101](#), [105](#), [109](#)
- [195] S. Ghosh, A. Bhagwat, C. Renshaw, S. Goh, A. Gaeta, and B. Kirby. Low-Light-Level Optical Interactions with Rubidium Vapor in a Photonic Band-Gap Fiber. *Physical Review Letters*, 97(2):023603, July 2006. [102](#)
- [196] T. Kawalec, M.J. Kaspruwicz, L. Jozefowski, A. Burchianti, C. Marinelli, and H.-G. Rubahn. Dynamics of laser-induced cesium atom desorption from porous glass. *Chemical Physics Letters*, 420(4-6):291–295, March 2006. [102](#)
- [197] S. Villalba, H. Failache, and A. Lezama. Light-induced atomic desorption and diffusion of Rb from porous alumina. *Physical Review A*, 81(3):032901, March 2010. [102](#), [107](#), [109](#)
- [198] C. Marinelli, A. Burchianti, A. Bogi, F. Della Valle, G. Bevilacqua, E. Mariotti, S. Veronesi, and L. Moi. Desorption of Rb and Cs from PDMS induced by non resonant light scattering. *The European Physical Journal D*, 37(3):319–325, November 2005. [102](#), [117](#)
- [199] A. Burchianti, A. Bogi, C. Marinelli, C. Maibohm, E. Mariotti, S. Sanguinetti, and L. Moi. Optical characterization and manipulation of alkali metal nanoparticles in porous silica. *The European Physical Journal D*, 49(2):201–210, August 2008. [107](#), [108](#), [112](#), [113](#), [117](#)
- [200] S. N. Atutov, V. Biancalana, and P. Bicchi. Light-induced diffusion and desorption of alkali metals in a siloxane film: Theory and experiment. *Physical Review A*, 60(6):4693–4700, 1999. [107](#), [109](#)
- [201] W. Hoheisel, M. Vollmer, and F. Träger. Desorption of metal atoms with laser light: Mechanistic studies. *Physical Review B*, 48(23), 1993. [107](#)
- [202] P. S. Donvalkar, S. Ramelow, S. Clemmen, and A. L. Gaeta. Continuous generation of rubidium vapor in hollow-core photonic bandgap fibers. *Optics Letters*, 40(22):5379, 2015. [113](#), [127](#)
- [203] S. D. Barrett, P. P. Rohde, and T. M. Stace. Scalable quantum computing with atomic ensembles. *New Journal of Physics*, 12, 2010. [118](#)

- [204] Y.-W. Cho, G. T. Campbell, J. L. Everett, J. Bernu, D. B. Higginbottom, M. T. Cao, J. Geng, N. P. Robins, P. K. Lam, and B. C. Buchler. Highly efficient optical quantum memory with long coherence time in cold atoms. *Optica*, 3(1):100, 2016. 121
- [205] D. Barredo, H. Labuhn, S. Ravets, T. Lahaye, A. Browaeys, and C. S. Adams. Coherent excitation transfer in a spin Chain of three rydberg atoms. *Physical Review Letters*, 114(11):1–5, 2015. 122
- [206] I. Novikova, A. V. Gorshkov, D. F. Phillips, A. S. Sørensen, M. D. Lukin, and R. L. Walsworth. Optimal control of light pulse storage and retrieval. *Physical Review Letters*, 98(24):1–4, 2007. 123
- [207] 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(6):1–5, 2017. 123
- [208] B. D. Bruner, H. Suchowski, N. V. Vitanov, and Y. Silberberg. Strong-field spatiotemporal ultrafast coherent control in three-level atoms. *Physical Review A - Atomic, Molecular, and Optical Physics*, 81(6):1–5, 2010. 123
- [209] S. Chakrabarti, A. Pradhan, B. Ray, and P. N. Ghosh. Velocity selective optical pumping effects and electromagnetically induced transparency for D2 transitions in rubidium. *J. Phys. B*, 38:4321–4327, 2005. 124
- [210] K. A.G. Fisher, D. G. England, J.-P. W. MacLean, P. J. Bustard, K. Heshami, K. J. Resch, and B. J. Sussman. Storage of polarization-entangled THz-bandwidth photons in a diamond quantum memory. *Physical Review A*, 96(1):1–5, 2017. 125
- [211] J. Brendel, N. Gisin, W. Tittel, and H. Zbinden. Pulsed Energy-Time Entangled Twin-Photon Source for Quantum Communication. *Physical Review Letters*, 82(12):2594—2597, 1999. 126
- [212] C. K. Hong, Z. Y. Ou, and L. Mandel. Measurement of Subpicosecond Time Intervals between Two Photons by Interference. *Physical Review Letters*, 59(18):2044–2046, 1987. 126
- [213] F. Blatt, T. Halfmann, and T. Peters. One-dimensional ultracold medium of extreme optical depth. *Optics Letters*, 39(3):446, 2014. 127
- [214] A. Burchianti, A. Bogi, C. Marinelli, E. Mariotti, and L. Moi. Optical stabilization of Rb vapor density above thermal equilibrium. *Journal of Modern Optics*, 57(14-15):1305–1310, February 2010. 127
- [215] F. Benabid, P. J. Roberts, F. Couny, and P. S. Light. Light and gas confinement in hollow-core photonic crystal fibre based photonic microcells. *Journal of the European Optical Society: Rapid Publications*, 4(c):09004, January 2009. 128
- [216] L. Xiao, N. V. Wheeler, N. Healy, and A. C. Peacock. Integrated hollow-core fibers for nonlinear optofluidic applications. *Optics express*, 21(23):28751–7, 2013. 128
- [217] M. Triches, A. Brusch, and J. Hald. Portable optical frequency standard based on sealed gas-filled hollow-core fiber using a novel encapsulation technique. *Applied Physics B: Lasers and Optics*, 121(3):251–258, 2015. 128

- [218] R. Ritter, N. Gruhler, W. H P Pernice, H. Kübler, T. Pfau, and R. Löw. Coupling thermal atomic vapor to an integrated ring resonator. *New Journal of Physics*, 18(10), 2016. 128
- [219] G. Corrielli, A. Seri, M. Mazzera, R. Osellame, and H. de Riedmatten. Integrated Optical Memory Based on Laser-Written Waveguides. *Physical Review Applied*, 5(5):1–7, 2016. 128
- [220] L. Ko, M. W Noel, J. Lambert, R. Löw, H. Weimer, and J. Nipper. RF-dressed Rydberg atoms in hollow-core fibres. *J. Phys. B: At. Mol. Opt. Phys.*, 49, 2016. 128
- [221] E. Vliegen, S. Kadlecsek, L. W. Anderson, and T. G. Walker. Density Measurements for Optically Thick Alkali Vapors Using the Faraday Effect. *Nucl. Instrum. & Methods A*, 460:444, 2001. 128
- [222] 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(15):155004, 2008. 148
- [223] J. H.D. Munns, C. Qiu, P. M. Ledingham, I. A. Walmsley, J. Nunn, and D. J. Saunders. In situ characterization of an optically thick atom-filled cavity. *Physical Review A*, 93(1):1–8, 2016. 149
- [224] A. K. Ekert and P. L. Knight. Correlations and squeezing of two-mode oscillations. *American Journal of Physics*, 57(8):692–697, 1989. 155

Appendix A. Optical depth

The optical depth of a medium characterises the amount of light that will be absorbed as it passes through the medium. According to Beer's law, the intensity of light after passing through a medium of optical depth d should reduce to $I = I_0 e^{-2d}$, where I_0 is the intensity of light before the medium. Operating on narrowband atomic resonances, one has to in principle treat d as a function of frequency ω [222]. And indeed, if one scans the frequency of a weak narrowband laser field over an atomic resonance, they will obtain a transmission spectrum of the atoms (see Fig. A.1 (c)).

The optical depth d can be related to the atomic number density n via the simplified expression

$$d = \frac{r^2 \omega_0 L}{2\gamma \epsilon_0 \hbar c} n, \quad (\text{A.1})$$

where r is the dipole matrix element for the transition we probe, ω_0 and γ are respectively the center frequency and the homogenous linewidth of that transition¹³, and L is the length of the medium [149].

More precisely, the optical depth is directly proportional to the product of the medium length and the imaginary part of the total atomic susceptibility χ , which is composed of the contributions from each atomic transition between the different hyperfine states involved in the Cs D2 line, i.e. $\chi = \sum_{i,f} \chi_{if}$. The susceptibility for the $|i\rangle \rightarrow |f\rangle$ transition is

¹³Since we do not use a buffer gas in our experiments, γ is only due to the spontaneous emission from the excited state.

$$\chi_{if} = c_{if}^2 \frac{n_i r_{if}^2}{\hbar \epsilon_0} s(\omega - \omega_{if}) \quad (\text{A.2})$$

where c_{if} is the transition strength factor and r_{if} is the dipole matrix element for that transition [223]. n_i is the number density of atoms in the $|i\rangle$ state and $s(\omega)$ is the resonance frequency distribution of the atoms around the $|i\rangle \rightarrow |f\rangle$ transition center frequency ω_{if} . For a warm atomic gas, we cannot neglect the contribution of the Doppler effect and thus $s(\omega)$ will have the shape of a convolution between the Lorentzian-shaped homogenous response and the Gaussian inhomogenous response of the atoms:

$$s(\omega) = \int_{-\infty}^{\infty} L(\omega - \omega') G(\omega') d\omega' \quad (\text{A.3})$$

where $L(\omega) = \frac{-\omega + i\gamma}{\omega^2 + \gamma^2}$, and $G(\omega) = \frac{1}{\sigma\sqrt{2\pi}} \exp(-(\frac{\omega}{\sigma})^2/2)$, σ being the Doppler width ¹⁴. The imaginary part of $s(\omega)$ gives the Voigt absorption profile.

The density of atoms in a vapour n , via the vapour pressure ¹⁵, can be related to the vapour's temperature using the empirical formula given in [103]. This allows one to estimate the actual cell temperature based on the measured optical depth.

In order to measure the cell temperature for the ORCA experiment, I used a Matlab code, initially developed by PhD student Joseph Munns, which fitted χ (including an additional scaling and slope term to compensate for the change in probe beam power over its frequency scan) to the measured transmission spectrum of the vapour cell. n_i , γ and σ were varied as free parameters in a least squares fit routine. The frequency axis of the transmission spectrum was calibrated using a simultaneously-registered saturated absorption spectrum in a different Cs cell, as shown in Fig. A.1 (a) and (b) (an explanation of saturated absorption can be found in Appendix B.1). The fit corresponding to the cell used in the ORCA experiment is shown in Fig. A.1 (c), and gives a cell temperature of $\sim 91^\circ\text{C}$.

For the LIAD-induced optical depth measurements, I used a less precise, but faster

¹⁴ $\sigma = \omega_{if} \sqrt{\frac{k_B T}{mc^2}}$ where m is atomic mass of caesium.

¹⁵ In an ideal gas, which is a good approximation for an atomic vapour at the observed densities, the number density $n = P_v/RT$, where P_v is the vapour pressure, V is the cell volume, R is the gas constant and T is the temperature.

and more robust, Matlab code written by myself, which used n_i and σ as fit parameters¹⁶, and fitted only one of the Cs D2 transition transmission spectra, i.e. the $6S_{1/2}(F=3) \rightarrow 6P_{3/2}(F=2,3,4)$ transitions. This was more appropriate for the large amounts of data that needed to be evaluated in this experiment. Since I expected to generate very high atomic densities using LIAD, the code also took into account self-broadening of γ ¹⁷ in the Voigt profiles [159]. Exemplary fits using that code for the cases of LIAD being off and on are shown in Fig. A.1 (d) and (e), respectively.

¹⁶The code would also fit the signal amplitude and background level, as these could change during the three hour long measurement.

¹⁷It is worth noting though that the observed increase in the homogenous linewidth was minimal (< 0.5 MHz).




Figure A.1: (a) Saturated absorption spectrum ([Appendix B. Laser locking](#)) of the Cs $6S_{1/2}(F=4) \rightarrow 6P_{3/2}$ transition, with marked hyperfine structure. (b) Same as (a), but for the Cs $6S_{1/2}(F=3) \rightarrow 6P_{3/2}$ transition. (c) Fitted absorption spectrum for the Cs D2 line in a bulk cell. (d) Fitted absorption spectrum for the Cs $6S_{1/2}(F=3) \rightarrow 6P_{3/2}$ transition in a hollow-core fibre without Light-Induced Absorption (LIAD). (e) Same as (d), but with LIAD on.

Appendix B. Laser locking




Figure B.1: Scheme of setup for locking the frequency of the (a) ECDL reference and (b) “signal” Ti:Sapph. HWP - half-wave plate; QWP - quarter-wave plate; PBS - polarising beamsplitter; SFP - scanning Fabry-Perót interferometer.

B.1 Locking the ECDL reference

In order to lock the frequency of our 852 nm continuous wave (CW) external-cavity diode laser (ECDL), I send a measured saturated absorption spectrum [193] in a room-temperature caesium reference cell into the ECDL controller, which internally produces an error signal that the ECDL locks to. It does so effectively by electronically differentiating the saturated absorption signal, yielding error signals at the peaks, as simulated in Fig. B.2. One of these error signals is then used for fast feedback on the ECDL diode current and slow feedback on the voltage on the piezoelectric element holding the ECDL cavity mirror.




Figure B.2: Saturated absorption signal (blue line) fed into the ECDL with simulated error signal (orange), similar to the one used for locking the laser.

To generate the saturated absorption spectroscopy signal, I use a few μW of ECDL power to probe the absorption in a room-temperature Cs reference cell, overlapped with a few mW of counter-propagating “pump”, as shown in Fig. B.1(a). The strong “pump” beam reduces the population of the atoms on the transition being investigated by the weak “probe” beam, and thus lowers the absorption of the latter. However, due to the counter-propagating beam configuration, pump and probe will experience opposite Doppler shifts from the moving atoms, and so in principle pump and probe will interact with different velocity classes of atoms for most laser frequencies. For certain laser frequencies however,

resonant with atoms in the zero velocity class (and a few “crossover” points¹⁸), the strong “pump” will interact with the atoms being probed, depleting their number and thus lowering absorption of the probe. This saturated (by the pump) absorption will result in narrow, Doppler-free peaks in the transmission spectrum, visible on Fig. B.2 or Fig. A.1.

B.2 Locking the signal Ti:Sapph

The center frequency of the 852 nm signal Ti:Sapph was locked with respect to the reference ECDL laser described in the previous section. A small pick-off from the Ti:Sapph was sent along with some of the locked reference CW light into a scanning Fabry-Pérot interferometer (SFP), as shown in Fig. B.1(b), which was connected to a computer running LabVIEW. First, one would manually tune the Ti:Sapph laser into resonance with the atoms, by observing fluorescence from a Cs reference cell on a camera. Then, a program written originally by former DPhil student Michael Sprague in LabVIEW was used to actively stabilise the desired frequency difference between the main peak of the Ti:Sapph spectrum and the CW laser, by calculating this difference (as measured by the SFP) and feeding back on the voltage on the Ti:Sapph intra-cavity Gires-Tournois Interferometer (GTI)¹⁹ [193]. This scheme stabilised the Ti:Sapph frequency to within approximately 100 MHz, at Hz locking rate.

B.3 Locking the control Ti:Sapph

I locked the 917 nm control Ti:Sapph center frequency by sending a small pick off into a commercial wavelength meter (High Finesse WS6-600) with a built-in electronic feedback system. Similar to the signal Ti:Sapph, I feedback on the voltage on the GTI inside the Ti:Sapph cavity. The accuracy of determining the center frequency in this system was 100 MHz, which was also the accuracy of the lock. The feedback speed was on the order of kHz, i.e. the measurement rate of the wavelength meter.

¹⁸These appear when transition $|g\rangle \rightarrow |e1\rangle$ in the $v_s - v$ velocity class is pumped, but probed by a neighbouring transition $|g\rangle \rightarrow |e2\rangle$ in the $v_s + v$ class. Here, v_s is the average speed of the ensemble, $|g\rangle$ is the ground state of the investigated transitions (the population of which is being measured by absorption of the probe), and $|e1, 2\rangle$ are neighbouring excited states of said transitions.

¹⁹By adjusting the voltage to a piezoelectric driver that controls the GTI mirror spacing, one can modify the latter’s dispersion, and thus the optical frequency at which the Ti:Sapph is encouraged to mode-lock.

Appendix C. $g^{(2)}$

The second-order auto-correlation function $g^{(2)}$ (as defined in section 4.1.2), which quantifies the photon number statistics of an optical field, is a good method to distinguish between different quantum and classical states of light. In terms of normally ordered photon creation $\hat{a}^\dagger$ and annihilation $\hat{a}$ operators, it can be defined as:

$$g^{(2)} = \frac{\langle (\hat{a}^\dagger)^2 (\hat{a})^2 \rangle}{\langle \hat{a}^\dagger \hat{a} \rangle^2}, \quad (\text{C.1})$$

which can also be related to the mean photon number $\langle n \rangle$ ²⁰ and variance V_n ²¹ by:

$$\frac{V_n - \langle n \rangle}{\langle n \rangle} = \langle n \rangle (g^{(2)} - 1), \quad (\text{C.2})$$

also known as Mandel's Q parameter [224]. It can be easily seen that when the photon number distribution is sub-Poissonian ($V_n < \langle n \rangle$), which is a signature of anti-bunching and thus quantum light²², then the Q parameter is negative and $g^{(2)} < 1$.

For example, for the quantum n-photon Fock state, i.e. a photon number state $|n\rangle$ ²³:

$$g^{(2)} = \frac{\langle n | \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} | n \rangle}{\langle n | \hat{a}^\dagger \hat{a} | n \rangle^2} = \frac{n(n-1)}{n^2} = 1 - \frac{1}{n} < 1. \quad (\text{C.3})$$

Specifically, the single photon state $|1\rangle$ will have a $g^{(2)} = 0$. Higher photon number contributions will necessarily lead to the $g^{(2)}$ increasing, as can be expected from Eq. C.3.

²⁰ $\langle n \rangle = \langle \hat{a}^\dagger \hat{a} \rangle$

²¹ $V_n = \langle n^2 \rangle - \langle n \rangle^2 = \langle (\hat{a}^\dagger)^2 (\hat{a})^2 \rangle + \langle \hat{a}^\dagger \hat{a} \rangle - \langle \hat{a}^\dagger \hat{a} \rangle^2$

²² Sub-Poissonian statistics violate the Cauchy-Schwartz inequality for classical distributions [149].

²³ $\hat{a}|n\rangle = \sqrt{n}|n-1\rangle$

Appendix D. Polarisation

In the Raman Λ -type memory, where storage occurs between two spin states mediated by an excited state manifold, there exist specific polarisation selection rules which impose that the signal and control must be orthogonally polarised for the memory to operate [23]. In the current implementation of the ORCA protocol, we instead store light in the hyperfine manifold of a doubly-excited state (as shown in Fig. 3.2, section 3.1.1). In fact, as discussed in section 3.2.2, we actually store the optical field as a coherent superposition of atomic coherences between the ground state and the different hyperfine levels composing the doubly-excited state manifold. The mapping of light to each of these coherences will utilise different polarisation selection rules, preventing the formation of any straightforward requirements on the polarisation of the signal and control fields, and making the polarisation properties of the memory quite complicated altogether.

For example, Fig. D.1 shows how, for a given horizontal signal polarisation, the ORCA storage efficiency changes for 6 different control field polarisations. The storage efficiency varies between 61% for a left-circular polarised control field and 83% for anti-diagonal polarised light.

Because in all the experiments described in this thesis our detection scheme was composed of a polarising beamsplitter (PBS) with a detector in each output, our detection was essentially polarisation sensitive, giving us a bit more information about the polarisation properties of ORCA. For example, Fig. D.2 shows a polarisation-sensitive version of the memory lifetime measurement described in section 3.2.2. As one can see, the memory efficiency is quite different in the two polarisation modes defined by the PBS. Furthermore, visible is weak oscillation in the memory lifetime curve with time, opposite in phase for




Figure D.1: Storage efficiency versus control field polarisation for horizontally polarised signal. The distance from the center of the sphere (of radius 1) is the experimentally determined storage efficiency for the given polarisation.

the two polarisation modes. This may be due to the beating of the different components of the superposition of stored coherences, different for the two polarisation modes, as the coefficients in these superpositions will differ for different optical field polarisations. This also means that the resultant polarisation of the retrieved signal may vary with storage time.

Naively, the polarisation effects described above preclude the use of the ORCA memory for storing polarisation qubits, as we have no guarantee the polarisation of the signal is preserved in the storage process. However, it may be possible to tune the polarisation of the control field such that there exists a polarisation eigenbase for the signal. It is also possible to prepare (by optical pumping) the memory ensemble in the highest (lowest) magnetic sub-level and then use σ^+ (σ^-) -polarised signal and control fields to create an atomic coherence only between specific magnetic sub-levels [18]. This would preserve the polarisation of the signal, thus enabling storage of polarisation qubits in a dual-rail configuration [86].




Figure D.2: Polarisation-sensitive memory lifetime curve. The measured total memory efficiency versus storage time for the two different outputs of the PBS composing the Hanbury-Brown Twiss detection setup. Errorbars are smaller than the data markers.