A Cavity-Enhanced Room-Temperature Broadband Raman Memory

D. J. Saunders[1,†], J. H. D. Munns[1,2], T. F. M. Champion[1], C. Qiu[1,3],

K. T. Kaczmarek[1], E. Poem[1], P. M. Ledingham[1], I. A. Walmsley[1], J. Nunn[1]

[1]*Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom*

[2] *QOLS, Blackett Laboratory, Imperial College London, London SW7 2BW, UK*

[3]*Department of Physics, Quantum Institute for Light and Atoms, State Key Laboratory of Precision Spectroscopy,*

East China Normal University, Shanghai 200062, People's Republic of China

(Dated: February 24, 2022)

Broadband quantum memories hold great promise as multiplexing elements in future photonic quantum information protocols. Alkali vapour Raman memories combine high-bandwidth storage, on-demand read-out, and operation at room temperature without collisional fluorescence noise. However, previous implementations have required large control pulse energies and suffered from four-wave mixing noise. Here we present a Raman memory where the storage interaction is enhanced by a low-finesse birefringent cavity tuned into simultaneous resonance with the signal and control fields, dramatically reducing the energy required to drive the memory. By engineering anti-resonance for the anti-Stokes field, we also suppress the four-wave mixing noise and report the lowest unconditional noise floor yet achieved in a Raman-type warm vapour memory, $(15 \pm 2) \times 10^{-3}$ photons per pulse, with a total efficiency of $(9.5 \pm 0.5)\%$.

Quantum information technologies such as quantum key distribution and random number generators are beginning to transition into the commercial sphere, where key requirements are the ability to function at high speed, and in non-laboratory settings. The high carrier frequency of optical signals enables photonic quantum devices to operate noise-free at room temperature whilst offering GHz-THz operational bandwidths. However, direct photon-photon interactions are prohibitively weak, which has held back the development of photonic quantum processors. One solution to this problem has been to use probabilistic measurement-induced non-linearities [1], but the probability of success decreases exponentially with system size, limiting the state of the art to < 10 photons [2]. Further scaling photonic devices requires a multiplexing strategy. Quantum memories capable of storing photons and releasing them on-demand provide the ability to temporally multiplex a repeat-until-success architecture to achieve a freely scalable photonic quantum information platform operable at room temperature.

There are many types of quantum memory for light under development: electromagnetically induced transparency [3], the full atomic-frequency comb protocol [4, 5], gradient-echo memories [6, 7], and the far-off-resonant Raman memory [8, 9]. Each of these protocols have advantages and challenges, see Ref. [10] for a recent review. A helpful metric for temporal multiplexing is the time-bandwidth product $B = \tau\delta$, with δ the acceptance bandwidth and τ the memory lifetime. B is the maximum number of time-bins over which a memory can synchronise an input signal. Time-bandwidth products of $B > 1000$ enable a dramatic enhancement in the multiphoton rate from parametric photon sources [11], as required to utilise multiphoton interference for computational speed-up (e.g. boson sampling [12]). Large time-bandwidth products have been achieved with cold atom memories [13, 14], cryogenic rare earth ion memories [15] and room-temperature Raman memories [8, 16]. In this paper we present an implementation of an alkali-vapour Raman memory, operating at 75°C, and show the lowest-noise floor demonstrated in such a room-temperature memory to date.

In a Raman memory, an intense control field induces a two-photon absorption feature far-detuned from any atomic resonance. The absorption linewidth of the field-dressed atoms can be made broadband with a pulsed control [16], enabling a large time-bandwidth product even with $\sim \mu$s coherence times in a room-temperature vapour [8], and allowing direct interfacing with pulsed heralded photon sources [17, 18]. Operation with temporally short and far-detuned pulses also removes contamination of the retrieved signals by collision-induced fluorescence [19]. However, large control pulse energies are needed to drive the storage interaction far from resonance, and the control can also drive unwanted four-wave mixing, introducing noise which cannot be filtered out either spectrally or temporally. Four-wave mixing noise has emerged as the last remaining roadblock to the development of efficient Λ-type room-temperature quantum memories [17, 20–23]. We solve this problem by introducing a new *cavity enhanced* Raman memory protocol. We place the atoms inside a low-finesse optical cavity that both enhances the strength of the Raman interaction, reducing the power requirements on the control field, and also suppresses four-wave mixing.

To describe how the cavity enhances the memory, we first introduce the standard Raman protocol, which we have developed in caesium vapour at ~ 70°C [8, 17, 25–27], using the Λ-system of the Cs D2 line comprising the hyperfine ground states $6S_{1/2}$, $F = 3$ ($|3\rangle$), $F = 4$ ($|1\rangle$), with a hyper-fine splitting of $\Delta_{\mathrm{HF}} = 9.2$ GHz and the $6P_{3/2}$ excited state manifold ($|2\rangle$), whose hy-
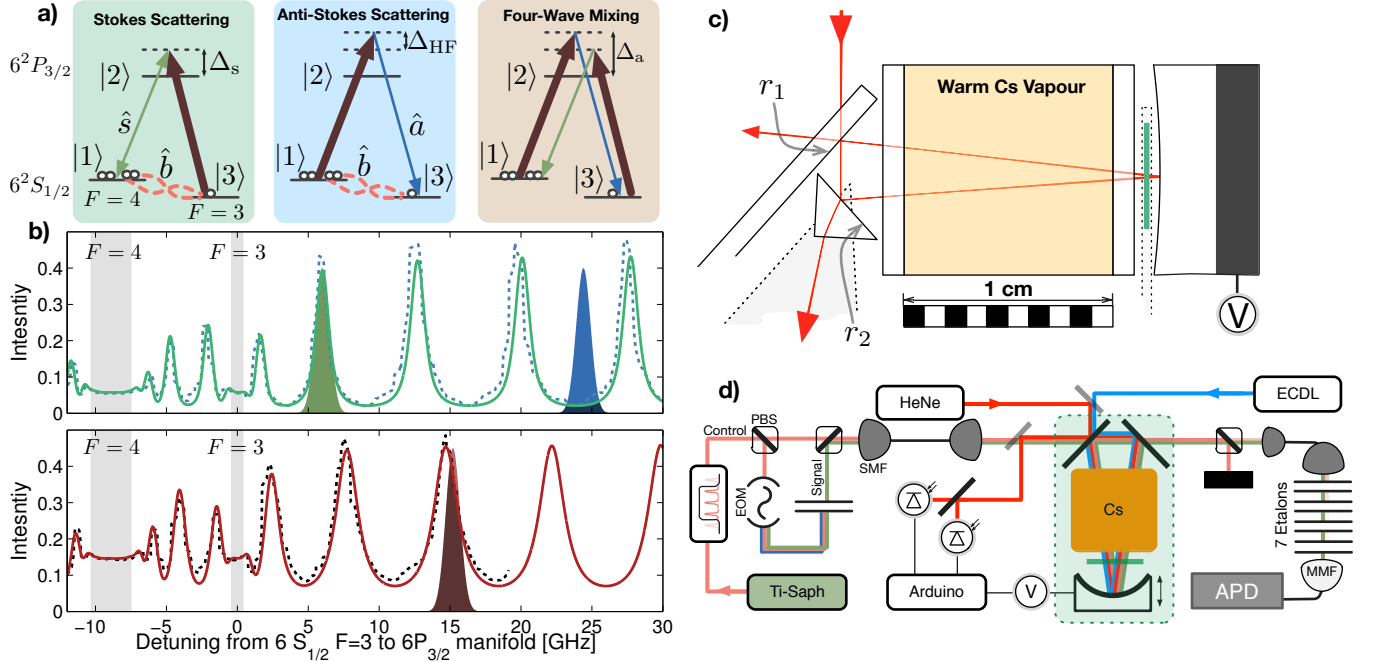
FIG. 1: (a) Energy level diagrams showing the D2 Λ level scheme in Cs. There are two major interactions: one, the Stokes scattering — involving the signal field $\hat{s}$ (green arrow), control field (red arrow), and spin-wave $\hat{b}$ (red loop), this is the desired memory interaction; two, anti-stokes scattering — involving the control field, anti-stokes mode $\hat{a}$ (blue arrow), and spin-wave $\hat{b}$; the third diagram is the effective noise-process, anti-stokes scattering followed by stokes-scattering, a four-wave-mixing interaction. (b) Measured frequency response of the atom-filled cavity system. The top is for the signal and anti-stokes polarisation, and the bottom is for the control polarisation. The dashed lines are measured responses (see supplementary material) and the solid lines are our theory model using measured values of the cavity parameters [24]. The grey solid areas represent the absorption in the cavity, and are an indication of our spin-polarisation, with $\sim 80\%$ in the ideal $F = 4$ state. The green pulse shows the location of the signal, the blue the anti-stokes, and the red the control. (c) Scale diagram of the cavity Raman quantum memory. We operate the ring-cavity in transmission, where in-coupling and out-coupling mirrors have amplitude reflectivities r_1 and r_2 respectively. The curved mirror has $R = 40$mm, and the small green rectangle represents the 200μm-thick LiNbO$_3$ birefringent crystal which we tune with temperature to reach the desired cavity conditions. (d) Experimental schematic, for more details see the main text and the supplementary material.

perfine structure is unresolved, see Fig. 1 (a). We initialise the memory by optically pumping the atoms into $|1\rangle$. Light storage and subsequent retrieval are mediated by the application of bright control pulses with bandwidth δ, time-dependent Rabi frequency $\Omega(t)$ and detuning $\Delta_{\rm s}$ from resonance with the $|3\rangle \leftrightarrow |2\rangle$ transition. This adiabatically couples the signal mode $\hat{s}$ to a delocalised excitation of the storage state $|3\rangle$ described by the spin-wave mode $\hat{b}$, according to the beam-splitter-like Hamiltonian $\mathcal{H}_{\rm BS} \propto C_{\rm s}\hat{s}\hat{b}^\dagger + {\rm h.c.}$, where the Stokes coupling parameter is given by $C_{\rm s} = \sqrt{Wd\gamma}/\Delta_{\rm s}$, with $W = \int |\Omega(t)|^2 {\rm d}t$ proportional to the control pulse energy $\mathcal{E}$ and d the resonant optical depth of the $|1\rangle \leftrightarrow |2\rangle$ transition, with linewidth γ [28]. The origin of the unwanted four-wave mixing noise is the coupling of the control field to the ground-excited transition $|1\rangle \leftrightarrow |2\rangle$ with detuning $\Delta_{\rm a} = \Delta_{\rm s} + \Delta_{\rm HF}$, on which it drives spontaneous anti-Stokes scattering in mode $\hat{a}$, producing anti-Stokes photons and spin-wave excitations in pairs [29]. The interaction is accordingly described by a two-mode-squeezing Hamiltonian $\mathcal{H}_{\rm TMS} \propto C_{\rm a}\hat{a}\hat{b} + {\rm h.c.}$, with the anti-Stokes

coupling given by $C_{\rm a} = C_{\rm s}\Delta_{\rm s}/\Delta_{\rm a}$. If $\Delta_{\rm s} \gtrsim \Delta_{\rm HF}$, the interactions described by $\mathcal{H}_{\rm BS}$, $\mathcal{H}_{\rm TMS}$ are of comparable strength, and the probability of spontaneously generating a spin-wave excitation approaches that of mapping an incident signal photon into the spin-wave. The retrieved fields are then contaminated with thermal noise (half a two-mode-squeezed state), which destroys the non-classical statistics of stored single photons [17].

One way to prevent four-wave mixing is to use polarisation selection rules such that the control field only couples to the $|3\rangle - |2\rangle$ transition. However, for alkali atoms this is not possible due to destructive interference between interaction pathways associated with different intermediate excited states $|2\rangle$ [22]. Four-wave mixing can also be suppressed in dispersive media where it is poorly phasematched [9, 18], but the Stokes shift $\Delta_{\rm HF}$ in alkali vapours is not large enough to introduce a significant phase mismatch. A third approach towards suppressing four-wave mixing is the creation of a second Raman absorption feature at the anti-stokes frequency [30]. Whilst promising, this method requires the use of

an atomic species with three ground states, and therefore cannot be implemented in ^{133}Cs.

The cavity-enhanced Raman memory instead suppresses $\mathcal{H}_{\mathrm{TMS}}$ relative to $\mathcal{H}_{\mathrm{BS}}$ by engineering the density of scattering states so that anti-Stokes scattering into the fundamental cavity mode is forbidden by destructive interference. That is, we arrange for the signal (Stokes) field to coincide with a cavity resonance, while the anti-Stokes field is anti-resonant with the cavity. Setting the free-spectral range of the cavity accordingly, with $\mathrm{FSR}_m = 4\Delta_{\mathrm{HF}}/(2m+1)$, for integer m, and a roundtrip optical path length $l_m = c/\mathrm{FSR}_m$, the spontaneous generation of spurious spin wave excitations by anti-Stokes scattering via $\mathcal{H}_{\mathrm{TMS}}$ is then reduced relative to $\mathcal{H}_{\mathrm{BS}}$ by the ratio of intra-cavity field amplitudes, parameterised by the noise suppression factor $x = (1 - \mu_{\mathrm{s}})/(1 + \mu_{\mathrm{a}}) \approx 0.24$, where $\mu_{\mathrm{s,a,\Omega}}$ denotes the cavity-roundtrip amplitude losses for the intra-cavity fields for the Stokes ($\mu_{\mathrm{s}} = 0.6$), anti-Stokes ($\mu_{\mathrm{a}} = 0.6$) and control ($\mu_{\Omega} = 0.4$) frequencies. Much greater noise suppression could be achieved with a moderate increase in the cavity quality factor.

With a cavity around the atoms, the acceptance bandwidth of the memory is now limited by the cavity resonance linewidth, so that a broadband memory requires a low-finesse cavity. However, unlike in single-atom cavity QED [31], where very high cavity quality factors are required to achieve the strong coupling regime and suppress spontaneous emission into any direction, here we require only the suppression of on-axis spontaneous anti-Stokes scattering, since off-axis scattering couples to momentum-orthogonal spin wave modes that are not phasematched [23] and do not contribute noise. Cavity-enhanced ensemble Raman storage is therefore technically less demanding than cavity-QED-based storage, requiring only low- to moderate-finesse, which remains compatible with broadband operation ($\delta \sim$ GHz).

The broadest acceptance bandwidth is obtained for a zero-order cavity with $\mathrm{FSR}_0 = 36.8$ GHz, but to achieve $l_0 = 8.1$ mm would require a monolithic construction. For our demonstration, we instead constructed a second-order cavity with $\mathrm{FSR}_2 = 7.36$ GHz and $l_2 = 40.8$ mm. We chose a ring geometry instead of a linear one so as not to produce a standing wave inside the cavity, since atoms diffusing into the field nodes during the storage time would not interact at retrieval. To in-couple the control field, which is linearly polarised orthogonal to the signal and anti-Stokes modes, we tune the cavity birefringence by heating a thin sliver of LiNbO$_3$ (see Fig. 1) to achieve simultaneous resonance for the signal and the control fields. The cavity length is stabilised with a Hänsch-Couillaud lock [32]: we inject a HeNe laser into the cavity and feed back on the polarisation rotation of the beam by adjusting the cavity length with a piezo-electric actuator (Fig. 1). An Arduino microcontroller (Due) provides both fast (kHz) and slow (mHz) feedback (see Supplementary Material).

Besides suppressing FWM noise, the double resonance configuration for the control and signal fields reduces the control pulse energy required to drive the memory by a factor $\sim \mathcal{F}_{\mathrm{s}}^2 \mathcal{F}_{\Omega}^2 / \pi^4$ compared to an equivalent free-space memory, where $\mathcal{F}_{\mathrm{s,\Omega}}$ denotes the cavity finesse for the signal and control modes. In our experiment with $\mathcal{F}_{\mathrm{s}} \sim 7$ and $\mathcal{F}_{\Omega} \sim 4$. This combined with sleigh imperfect resonance of the control field (Fig. 1 (b)), we obtained a reduction in energy by ~ 10; a moderate increase in cavity quality could achieve extremely low operating power.
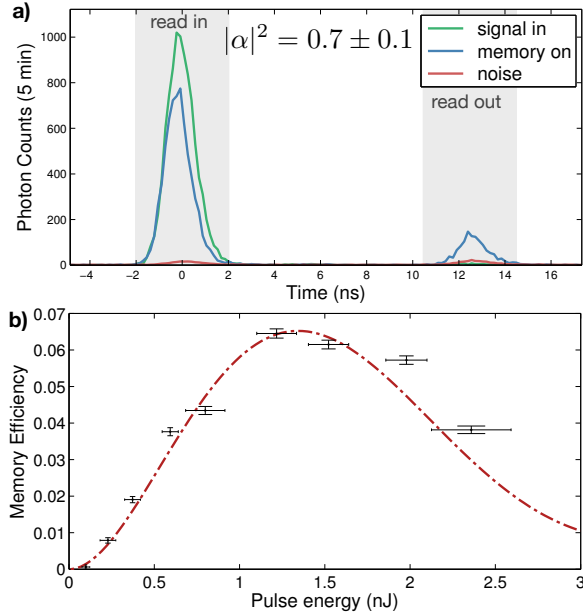

FIG. 2: (a) Arrival time histograms for the signal field emerging from the cavity with the control pulses blocked (*signal in*), unblocked (*memory on*) and with the input signal field blocked (*noise*). This is measured for a coherent state amplitude of $|\alpha^2| = 0.7 \pm 0.1$. (b) Measured (points) and predicted (line) memory efficiency, as a function of the intra-cavity control pulse energy.

As can be seen from Fig. 1(b), absorption and dispersion around the atomic resonances significantly modify the cavity mode structure [24]. The Cs vapour pressure depends exponentially on temperature, so to maintain the resonance condition required for the memory we developed active temperature stabilisation to within $0.1°$C, see supplementary material for details.

The experimental setup for the memory is shown in Fig. 1(d). The signal and control fields are derived from a modelocked Ti:Sapph oscillator producing pulses with 320 ps duration ($\delta = 1.2$GHz); details can be found in [17]. The signal and control pulses are coupled into the same single-mode fibre and directed towards the cavity memory. After the memory we block the control pulses with polarisation and spectral filtering, with a rejection of 120 dB and a transmission for the signal of 1%. We detect the transmitted and retrieved signal photons with a Geiger-mode avalanche photodiode (Perkin-
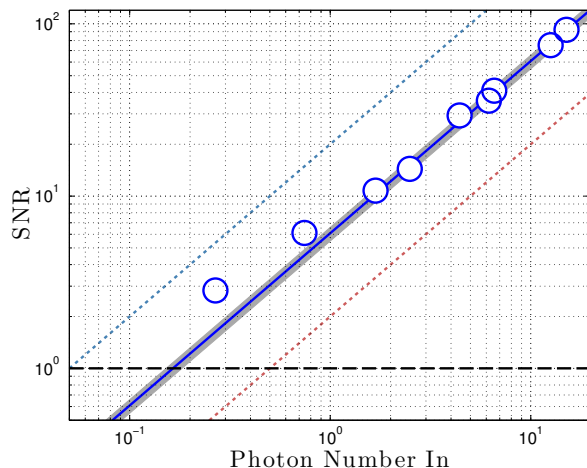
FIG. 3: Measurement of μ_1. To measure μ_1 we operate the memory with a control pulse energy of ~ 1.4 nJ, corresponding to our optimal SNR condition. We then measure the signal-to-noise ratio for various input coherent state amplitudes. Where the fits cross the horizontal-black dashed line (SNR=1) is the value of μ_1; with smaller μ_1 desirable. The blue circles are our measured values; error bars are smaller than marker dimension. The lower-red dashed line represents the state-of-the-art using cavity-free Raman memory experiments [17], with a $\mu_1' \approx 0.5$. The upper-blue dashed line is the expected performance from our theoretical model, with $\mu_1^{\text{Ideal}} = 0.005$. We do not reach this ideal noise-floor due to experimental imperfections not included the model (see supplementary material). After fitting our data with $y = x/\mu_1$, we extract $\mu_1 = 0.17 \pm 0.02$. The error is calculated using a Monte-Carlo simulation.

Elmer SPCM-AQRH) and time-to-digital convertor with 81ps resolution (qutools quTAU). Fig 2(a) shows storage and retrieval with consecutive control pulses separated by 12.5 ns. The input signal average photon number is 0.7 ± 0.1 and the memory efficiency is $(9.5 \pm 0.5)\%$, while the noise floor is $\langle n^{\text{noise}} \rangle = 0.015 \pm 0.002$ photons per control pulse in the read-out time bin, the lowest noise floor measured in warm vapour Raman memories to date. We also measure the effect of memory efficiency with control pulse energy, Fig 2(b). We observe a maximum efficiency at $\mathcal{E} \sim 1.5$nJ, before the memory efficiency is decreased due to the dynamical AC-stark effect. Nevertheless, the measured efficiency agrees with our new theoretical model, see supplementary material and ref. [33]. This effect can be compensated by using appropriate pulse shaping [20, 25, 34]. We also measured the memory lifetime to be $\tau = 95 \pm 7$ns, without any magnetic shielding (see supplementary information). This agrees with our previous measurements without shielding and we expect to easily extend this to $> \mu s$ with magnetic field shielding in the future [8].

The noise floor in our experiment contains contributions from four-wave mixing and also from spontaneous emission. To compare with the theory, we fit our measurements and extract the FWM $\mathcal{E}^2$-scaling component of the noise floor to find 0.006 ± 0.003, which agrees well with the theoretical prediction of 0.005 (Supplementary Information). Furthermore, a convenient metric to determine if we have suppressed the noise is the μ_1 parameter [4], where $\mu_1 = \langle n^{\text{noise}} \rangle / \eta^{\text{tot}}$; a fair metric to compare different devices because it includes both memory efficiency and noise floor. To determine this accurately, we vary the coherent state amplitude and determine the SNR for each, shown in Fig 3. By fitting this data, we determine a raw μ_1 of 0.17 ± 0.02, a considerable improvement compared to free-space Raman memories, for which $\mu_1' \sim 0.5$ [17]. This is a clear demonstration of noise suppression in a far-off resonant broadband Raman memory for the first time.

In order to take full advantage of the cavity-induced FWM noise suppression, it will be necessary to remove the remaining sources of noise. This can be achieved by improving the optical pumping efficiency, and suppression of counts arising from the pumping laser. The task at hand therefore reduces to a readily solvable engineering challenge. Once achieved, we can predict the expected performance of a future cavity Raman memory with heralded single photon inputs. Even at the current moderate memory efficiencies, and using SPDC sources which have previously been shown to be compatible with broadband Raman memories [17], interfacing them with our new experiment should enable non-classical revival of single-photons [33]. A realistic route to increasing the memory efficiency beyond 10% is to compensate for the AC-stark shift using pulse shaping, increasing the caesium number density, or increasing the finesse. Therefore, our new cavity enhanced Raman memory has laid the ground work for a true room-temperature quantum memory for broadband heralded single photons in the near future.

† Electronic Address: dylan.saunders@physics.ox.ac.uk

[1] J. L. O'Brien, A. Furusawa, and J. Vuckovic, Nat Photon **3**, 687 (2009).

[2] X.-C. Yao, T.-X. Wang, P. Xu, H. Lu, G.-S. Pan, X.-H. Bao, C.-Z. Peng, C.-Y. Lu, Y.-A. Chen, and J.-W. Pan, Nature Photonics **6**, 225 (2012).

[3] Y.-H. Chen, M.-J. Lee, I.-C. Wang, S. Du, Y.-F. Chen, Y.-C. Chen, and A. Y. Ite, Phys. Rev. Lett. **110**, 083601 (2013).

[4] M. Gündoğan, P. M. Ledingham, K. Kutluer, M. Mazzera, and H. de Riedmatten, Phys. Rev. Lett. **114**, 230501 (2015).

[5] P. Jobez, C. Laplane, N. Timoney, N. Gisin, A. Ferrier, P. Goldner, and M. Afzelius, Phys. Rev. Lett. **114**, 230502 (2015).

[6] M. P. Hedges, J. J. Longdell, Y. Li, and M. J. Sellars, Nature **465**, 1052 (2010).

[7] M. Hosseini, G. Campbell, B. M. Sparkes, P. K. Lam, and B. C. Buchler, Nat Phys **7**, 794 (2011).

[8] K. Reim, P. Michelberger, K. Lee, J. Nunn, N. Langford, and I. Walmsley, Phys. Rev. Lett. **107**, 053603 (2011).

[9] D. England, P. Bustard, J. Nunn, R. Lausten, and B. Sussman, Phys. Rev. Lett. **111**, 243601 (2013).

[10] Simon, C., Afzelius, M., Appel, J., Boyer de la Giroday, A., Dewhurst, S. J., Gisin, N., Hu, C. Y., Jelezko, F., KrÃll, S., MÃŒller, J. H., et al., Eur. Phys. J. D **58**, 1 (2010).

[11] 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, Phys. Rev. Lett. **110**, 133601 (2013).

[12] S. Aaronson and A. Arkhipov, in *Proceedings of the forty-third annual ACM symposium on Theory of computing* (ACM, 2011), pp. 333–342.

[13] X.-H. Bao, A. Reingruber, P. Dietrich, J. Rui, A. Dück, T. Strassel, L. Li, N.-L. Liu, B. Zhao, and J.-W. Pan, Nature Physics **8**, 517 (2012).

[14] Y. Dudin, L. Li, and A. Kuzmich, Phys. Rev. A **87**, 031801 (2013).

[15] G. Heinze, C. Hubrich, and T. Halfmann, Phys. Rev. Lett. **111**, 033601 (2013).

[16] P. J. Bustard, R. Lausten, D. G. England, and B. J. Sussman, Phys. Rev. Lett. **111**, 083901 (2013).

[17] 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, et al., New Journal of Physics **17**, 043006 (2015).

[18] D. G. England, K. A. Fisher, J.-P. W. MacLean, P. J. Bustard, R. Lausten, K. J. Resch, and B. J. Sussman, Phys. Rev. Lett. **114**, 053602 (2015).

[19] S. Manz, T. Fernholz, J. Schmiedmayer, and J.-W. Pan, Phys. Rev. A **75**, 040101 (2007).

[20] N. B. Phillips, A. V. Gorshkov, and I. Novikova, Phys. Rev. A **78**, 023801 (2008).

[21] N. Lauk, C. O'Brien, and M. Fleischhauer, Phys. Rev. A **88**, 013823 (2013).

[22] I. Vurgaftman and M. Bashkansky, Phys. Rev. A **87**, 063836 (2013).

[23] M. Dabrowski, R. Chrapkiewicz, and W. Wasilewski, Optics express **22**, 26076 (2014).

[24] J. Munns *et. al.*, In prep. (2015).

[25] K. Reim, J. Nunn, V. Lorenz, B. Sussman, K. Lee, N. Langford, D. Jaksch, and I. Walmsley, Nature Photonics **4**, 218 (2010).

[26] D. England, P. Michelberger, T. Champion, K. Reim, K. Lee, M. Sprague, X. Jin, N. Langford, W. Kolthammer, J. Nunn, et al., Journal of Physics B: Atomic, Molecular and Optical Physics **45**, 124008 (2012).

[27] M. Sprague, P. Michelberger, T. Champion, D. England, J. Nunn, X.-M. Jin, W. Kolthammer, A. Abdolvand, P. S. J. Russell, and I. Walmsley, Nature Photonics **8**, 287 (2014).

[28] J. Nunn, I. Walmsley, M. Raymer, K. Surmacz, F. Waldermann, Z. Wang, and D. Jaksch, Phys. Rev. A **75**, 011401 (2007).

[29] W. Wasilewski and M. G. Raymer, Phys. Rev. A **73**, 063816 (2006).

[30] G. Romanov, C. O'Brien, and I. Novikova, arXiv:1509.08951 (2015).

[31] H. P. Specht, C. Nölleke, A. Reiserer, M. Uphoff, E. Figueroa, S. Ritter, and G. Rempe, Nature **473**, 190 (2011).

[32] T. W. Hansch and B. Couillaud, Optics Communications **35**, 441 (1980).

[33] J. Nunn *et. al.*, In prep. (2015).

[34] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, Phys. Rev. A **76**, 033804 (2007).

[35] B. Beauregard, *http://playground.arduino.cc/code/pidlibrary*.

[36] A. Dantan, A. Bramati, and M. Pinard, Phys. Rev. A **71**, 043801 (2005).

Supplementary Material

Experimental set-up

The control and signal pulses are both derived from a mode-locked Tsunami titanium-sapphire (Ti:Saph) laser, producing an 80-MHz train of pulses with a centre wavelength of 852 nm and FWHM bandwidth of 1.2 GHz. The Ti:Sa frequency is actively stabilised at a detuning of $\Delta = 15.2$ GHz relative to the $6^2S_{1/2}(F = 3) \leftrightarrow 6^2P_{3/2}$ transition of a Cs reference cell. A Pockels cell (Quantum Technology QC) is used to select pulses for the storage and retrieval processes, setting an experimental repetition rate of 8 kHz. The signal field is generated with an electro-optic modulator (EOM) which is designed to shift the frequency of the modulated sidebands by an amount equal to the Cs ground-state hyperfine splitting Δ_{HF}. The red-detuned sideband, which satisfies the two-photon resonance condition with the control field, is frequency-selected using an air-spaced Fabry-Perot etalon. The orthogonally polarised signal and control pulses are temporally and spatially overlapped and coupled into a two-junction ring cavity in a co-propagating configuration. In addition, a narrowband Toptica DL pro external cavity diode laser (ECDL) is introduced into the cavity in a counter-propagating fashion in order to prepare the atoms in the initial ground state $|1\rangle$ via optical pumping.
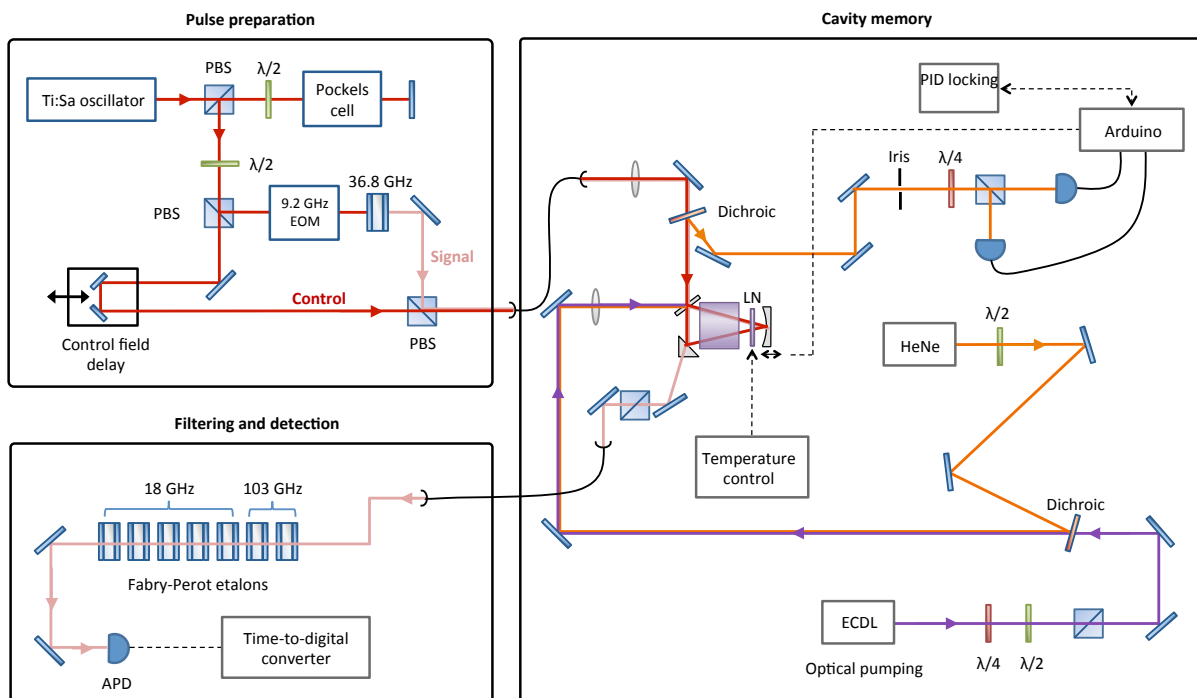

FIG. S1: Detailed schematic of the experiment. The experiment is separated into three parts: one, the pulse preparation; two, the cavity memory; and three, the filtering and detection. To prepare the pulses we pick pulses from our Ti:Sa oscillator and then split them on a PBS. From this PBS the two output paths are matched such that when they recombine on the second PBS they are overlapping temporally. We use an EOM to create the signal pulses, see main text for details. After this the signal and control pulses are coupled into the same fibre and sent to the cavity memory. On the cavity memory table we align all four fields, the pumping laser, the control and signal, and the locking laser to our ring-cavity. The output of the cavity is first filtered using a Glan-Taylor PBS before being coupled into single-mode-fibre and being sent to the filtering stage. Here the signal is passed through an effective total of 7 Fabry-Perot Etalons, with a control extinction of > 90dB before being coupled in multi-mode-fiber and coated on an APD and our time-tagging system. The total transmission of the measurement stage is $\sim 2\%$.

The cavity length is actively stabilised using a variant of the Hänsch-Couillaud method whereby an error signal is generated by monitoring the difference signal between orthogonal polarisation components of the HeNe locking laser transmitted through the cavity, arising from the relative phase accumulated on passing through the birefringent crystal, monitored by the 12-bit ADC of a microcontroller (Arduino Due) which then provides feedback to a piezo-electric ring actuator (Noliac) which controls the position of the concave cavity mirror along the axis of symmetry.

As a result of the periodicity of the error signal with cavity length being significantly different to the cavity response at the signal wavelength, it was necessary for the cavity length to remain locked to the correct set point solution of the error signal (corresponding to maximal signal coupling): in order to ensure this was the case, the feedback to the actuator consisted of two components: (1) a fast PID calculation, based on the Arduino library [35], with an output spanning the voltage range corresponding to a length change of the cavity associated with one period at the signal wavelength, and (2) a slow offset component which increments when the PID output approaches 90% of its range, spanning the whole voltage range of the actuator. Due to the finite precision of the DAC of the microcontroller, the two output components were output seperately and combined on a voltage adding circuit in the appropriate ratio to maximise the sensitivity of the output. This set up allows the lock to follow slower period drifts of the cavity whilst still remaining stable on the kHz timescale, and enabled the cavity to remain locked to within $\pm 50\,\mathrm{Hz}$ for a duration of hours.

The temperature of the Cs vapour cell is $84.3°\mathrm{C}$. This ensemble has both a large measured pressure broadening (FWHM = $80\,\mathrm{MHz}$), from 10 Torr of Ne buffer gas, and a measured Doppler broadening (FWHM = $375\,\mathrm{MHz}$). This makes both numerical modelling challenging [24]. This puts tight bounds on temperature stability, with changes of less than a degree causing large GHz size shifts in the cavity response; easily moving us away from memory operation. To solve this we developed a temperature control system to keep the temperature to within 0.1 of the target temperature. We monitored the temperature of the cell using thermistors and another Arduino Due, then using LabView we implemented a PID loop controlling the voltage to the vapour cells heating wires. The last step is to ensure the cavity is well aligned and has good fringe visibility — the key parameter in FWM suppression. To measure this we scan the cavity length using the piezo mirror, enabling us to measure $I_{\max}(f)$ and $I_{\min}(f)$; the maximum and minimum transmission intensities of the cavity for a given frequency, f. The cavity response in these measurements is affected by the linewidth of our probe laser ($\delta = 1.2\mathrm{GHz}$). We measure visibilities, $V = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}$, of $(86 \pm 2)\%$ for the Stokes frequency and polarisation, and $(71 \pm 5)\%$ for the Control frequency and polarisation, and $(86 \pm 5)\%\%$ for the anti-Stokes frequency and polarisation. The differences between these visibilities are ascribed to the change in mirror reflectivity for the control and signal, anti-Stokes polarisations, and a small amount of residual linear absorption at the signal frequency [24].

To ensure we are operating the memory in the optimal *triple resonance* condition — Stokes resonant, control resonant, anti-Stokes anti-resonant — we lock the cavity to the Stokes frequency and then scan our Ti-Saph laser from the D2-line to 30GHz blue detuned; limited by the $\sim 30\mathrm{GHz}$ mode-hop-free tuning range. We then measure the control polarisation, and tune the temperature of the birefringent element to make the control field also resonant. An example cavity frequency scan is shown in Figure 1b. Reaching triple resonance was made challenging because of dispersion induced by the Cs ensemble, as such we developed a full theoretical model, see [24].

The signal field transmitted through the cavity is coupled into a single-mode optical fibre and isolated from the control field using a series of polarisation and frequency filters. The arrival time of detection events relative to a trigger signal are recorded using a time-to-digital converter (qutools quTAU) with a time bin resolution of 81 ps.

Memory efficiency and memory lifetime

The memory efficiency can be evaluated from the rates of detection events $c_l^{\mathrm{in/out}}$ in the read-in and read-out time bins with $l \in \{sc, s, c\}$ denoting the combination of signal and control fields present. Note that the optical pumping laser, which is used to initiate the atomic ensemble prior to each storage event, is periodically switched off during the memory interaction to avoid depletion of the stored spin-wave excitation. The total memory efficiency for storage followed by retrieval is given by

$$\eta_{\mathrm{tot}} = \frac{c_{\mathrm{sc}}^{\mathrm{out}} - c_{\mathrm{c}}^{\mathrm{out}} - c_{\mathrm{s}}^{\mathrm{out}}}{c_{\mathrm{s}}^{\mathrm{in}}}, \tag{1}$$

Equation (1) assumes that the number of counts in the signal mode can be determined by subtracting the noise counts $c_{cd}^{\mathrm{in/out}}$ in each time bin. It is worth clarifying this assumption: The measured detection events consist of contributions from the signal (Stokes) mode, the anti-Stokes mode and the control mode. The latter contribution is the result of residual leakage of the control field through the filtering stage. We have implicitly assumed that the amount of control leakage is independent of the presence of the signal field, i.e. these leakage counts cancel in the expression $c_{\mathrm{scd}}^{\mathrm{out}} - c_{\mathrm{cd}}^{\mathrm{out}}$. The count rates $c_l^{\mathrm{in/out}}$ are found by integrating the pulse traces in Figure 2 (a) within the range indicated by the shaded areas. This gives a total memory efficiency of $\eta_{\mathrm{tot}} = 9.5 \pm 0.5\%$.

To measure the memory lifetime we vary the time between our read-in and read-out control pulses using our pulse picker, see [8] for details. The results are shown in Fig. S2.
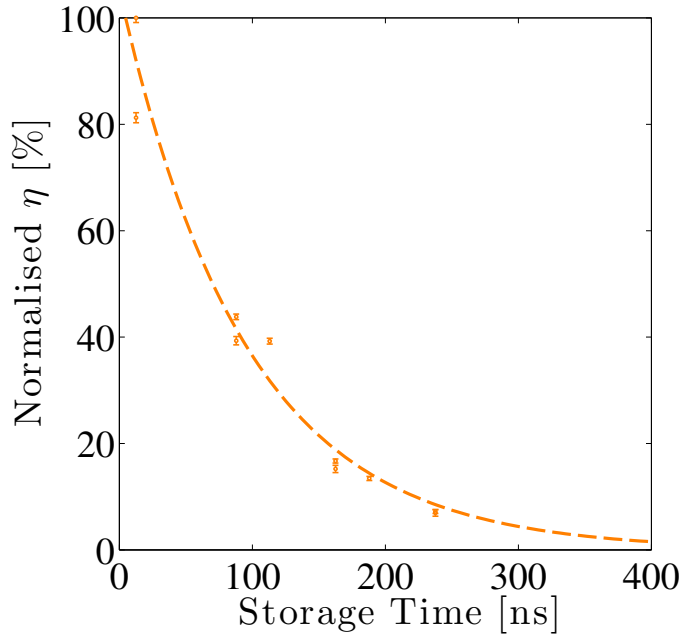
FIG. S2: We measured the memory lifetime by storing input pulses with average photon number ~ 10 and recording the average number of retrieved photons, for a range of storage times. We extract a memory lift time of 95 ± 7 ns from the exponential decay.

Theory of cavity Raman Interaction

The theory of cavity-enhanced Λ-type quantum memories has been treated in detail [34, 36]. We model the efficiency and noise in our Raman cavity memory by extending these analyses to include both Stokes and anti-Stokes fields [33]. Fig. 2(c) shows good agreement between the measured memory efficiencies and the prediction of our theoretical model, with no free parameters, given by

$$\eta_{\mathrm{tot}} = |\chi C_{\mathrm{s}} E \kappa|^2,$$

where $\chi = (1 - r^2)/(1 - \mu_{\mathrm{s}})$ is the amplitude transmission of the signal field with input- and output-coupling mirror amplitude reflectivities $r_1 = r_2 = r \approx 0.86$. $\kappa = (e^{\zeta} E)^{-1/2}(1 - e^{-f})/f$ is the input signal overlap with the cavity response, where $\zeta = -2\Re\{f\}$ and $E = (1 - e^{-\zeta})/\zeta$, with

$$f = -C_{\mathrm{s}}^2 - C_{\mathrm{a}}^2 x + \mathrm{i}\frac{W}{\Delta_{\mathrm{s}}} + \mathrm{i}\frac{W}{\Delta_{\mathrm{a}}}.$$

Since the intra-cavity fields pass through the atoms many times, the optical depth d is replaced by the cavity cooperativity $\mathcal{C} = rd/(1 - \mu_{\mathrm{s}}) \approx \mathcal{F}d/\pi$, where in our experiment the cavity finesse is $\mathcal{F} \approx 7$. Operationally, a measurement of the visibility $V = 90\%$ of the cavity fringes in Fig. 1 provides an estimate of the noise suppression via the relation $x^2 = (1 - V)/(1 + V)$. Our model then predicts the noise floor of the memory to be

$$N_{\mathrm{noise}} = |\chi C_{\mathrm{s}} C_{\mathrm{a}} x|^2 \times \frac{1 - e^{-\zeta} E}{\zeta}$$

photons per pulse. Using a measured beam-waist of $w = 150$ μm, we calculate $W[\mathrm{GHz}] \approx 110 \times \mathcal{E}[\mathrm{nJ}]$. We measure $V = 90\%$ and infer a resonant single-round-trip optical depth $d = 300$ with a pressure-broadened linewidth $\gamma = 250$ MHz [24].

We then obtain the prediction $N_{\mathrm{noise}} = 0.005$ photons per pulse for the maximum efficiency shown in Fig. 2(a).

Calculation of FWM Contribution

In order to isolate the contribution to the noise floor from FWM, we measured the scaling of the noise-floor as a function of the control pulse energy $\mathcal{E}$. The FWM contribution is expected to scale quadratically with $\mathcal{E}$, since
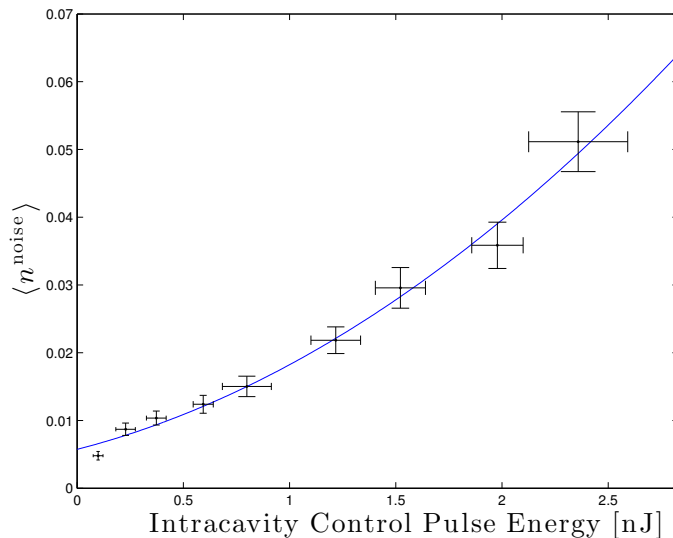
FIG. S3: Measured scaling of the noise-floor in our cavity memory.

two control pulse photons need to scatter to produce each FWM noise photon. The fit reveals the presence of a constant noise background and a component linear in $\mathcal{E}$. These two terms arise from control-field-independent noise, such as dark counts or fluorescence, and spontaneous Raman scattering from un-pumped population in the storage state, respectively. By extracting the quadratic term, we find a FWM contribution to the measured noise floor of 0.006 ± 0.003 for the control-field energy corresponding to our optimal SNR, which agrees with the prediction of our model (see above). With technical improvements to our optical pumping, our detection stage and our buffer gas (to further reduce collisional noise), we would be limited only by the FWM noise, which as we have shown can be removed with an appropriate cavity design.