Photon-echo rephasing of spontaneous emission from an ensemble of rare-earth ions

Sarah E. Beavan

January 2012

A thesis submitted for the degree
of Doctor of Philosophy of
The Australian National University

CHAPTER 6

Four-level photon echo

In this chapter a novel photon-echo in a four-level energy system is characterized. Here it is shown that this pulse sequence will allow for a photon echo to be detected down to a single-photon sensitivity level. This is the theoretical limit for a conventional photon-echo in a two-level system, however there are usually experimental limitations on reaching this limit, for instance due to non-uniformity in the rephasing pulse areas.

This photon echo sequence will be useful for rephasing spontaneous emission as in the RASE protocol (which is the topic of the next chapter). In this chapter, the pulse sequence will be characterized using an input pulse to generate ensemble coherence to be rephased.

The chapter outline is as follows: The first two sections discuss the motivation for moving to a four-level system. Following this, the rephasing process is theoretically considered, and the phase-matching condition is derived. The remainder of the chapter presents experimental results. This includes measurements of efficiency, coherence times, and echo-phase, and added noise. The the four-level echo concept was developed in conjunction with M. J. Sellars, and the phase-matching condition (presented in §6.3) was originally determined by P. Ledingham and J. J. Longdell.

6.1 Free induction decay (FID)

In a photon echo sequence, the role of the π-pulse is to reverse the effect of inhomogeneity in the transition frequency across an ensemble of atoms. In practice, most atoms distributed throughout the ensemble will experience a pulse area greater than or less than π, for instance due to variation in intensity across the beam profile.[1] An ensemble of atoms initially in the ground state will be perfectly inverted under the action of an ideal π-pulse. However following a realistic π-ish pulse, the Bloch vector describing the ensemble will have non-zero components in the $\hat{x}$-$\hat{y}$ plane of the Bloch-sphere (as discussed in §3.4.4), and this oscillatory component of the macroscopic polarization results in emission in the same

[1]Rabi frequency variation can also arise if the atoms have different oscillator strengths, or if the atoms are off-resonant, however all experiments in this thesis are performed on a sub-ensemble of ions of the same frequency subgroup, and the frequency width of this sub-ensemble is much less than the applied rephasing pulses, leaving spatial non-uniformity of the beam as the dominant effect for Rabi-frequency inhomogeneity.

spatial and spectral mode as the π-pulse. This emission during the free-induction decay (FID) will persist until the atoms dephase and the macroscopic polarization disappears. The FID decay time is T_2^* (see §3.4.1).

In the standard two-level echo (2LE) sequence, the echo is generated in the spatial and spectral mode defined by the π-pulse (the phase-matching condition is given by equation 3.38), which is then only distinguishable from the FID in the temporal domain. In the RASE protocol, it becomes particularly problematic to resolve the echo of one single photon amidst the many thousands of photons of FID emission. There are two strategies that could be applied to improve the quality of a π-pulse:

1. Engineer the beam to have a flat spatial distribution over the relevant ensemble of atoms, for instance use an aperture to select only the middle section of a gaussian beam. The disadvantage of this approach is that most of the available laser power is discarded.

2. Only select ions for the ensemble which see the same Rabi frequency. This can be achieved, for instance, by applying a 2π-pulse[2] followed by a delay longer than the excited state lifetime, and repeating this many times. As a result, atoms that see a pulse area other than 2π are eventually pumped via the excited state into a different (non-resonant) ground state, and removed from the experiment ensemble. This was demonstrated in reference [102], where after just 5 cycles, the variation in Rabi frequency across the remaining ensemble was $< 8\%$. The disadvantages of this technique, are the increase in the ensemble-preparation time, and that the number of potentially contributing ions is drastically reduced.

An alternative solution, which does not require perfect π-pulses, is to make the FID emission distinguishable from the echo.

6.2 Why four is better than two

By using a double-Λ system (four energy levels where all transitions are allowed) rather than a two-level system, the spectral overlap between echo and FID modes can be avoided altogether - it is possible to generate an echo on a transition that is not driven with any bright pulses.

This four-level echo (4LE) is depicted in figure 6.1 in the context of transitions in Pr dopant ions. The subset of levels $|2\rangle$, $|3\rangle$, $|4\rangle$ and $|5\rangle$ were chosen because they have similar transition strengths. The 4LE sequence begins with an input pulse generating coherence on the $|2\rangle \rightarrow |5\rangle$ transition (frequency ω_{25}), then after a delay time τ_a a π-pulse at ω_{35} transfers the coherence to a superposition between states $|2\rangle$ and $|3\rangle$. After a further delay τ_b a second π-pulse, at frequency ω_{24}, is applied. Then the coherence between states

[2] effectively a pulse with 'zero' net area

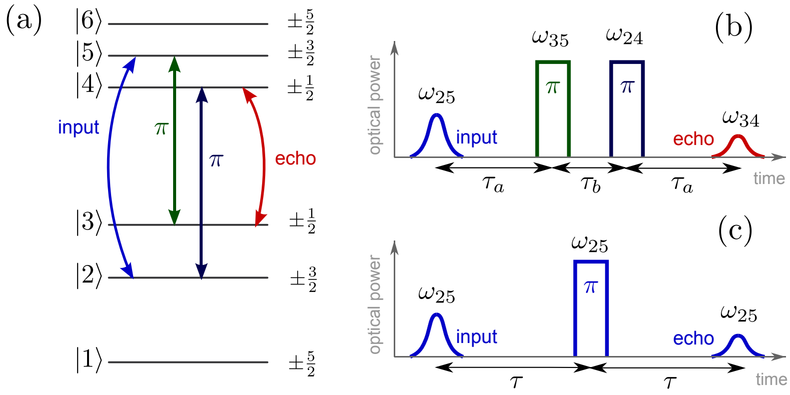

Figure 6.1: The four-level photon-echo sequence (4LE). (a) Reduced energy-level diagram of Pr:YSO with transitions of applied pulses for the 4LE sequence marked. (b) Pulse sequence for the 4LE, labelled with transition frequencies. The total delay time between input pulse and echo is $2\tau_a + \tau_b$. (c) Pulse sequence for the standard 2-level echo used for comparison measurements of echo efficiency and decoherence.

$|3\rangle$ and $|4\rangle$ evolves such that after further time τ_a the ensemble rephases and an echo is generated.

The key point to note is that the echo frequency (ω_{34}) is *different* from each of the three other input pulse frequencies, and thus is spectrally resolvable from FID emission. To achieve this requires a minimum of four energy levels. Although photon echoes can also be generated in three-level systems (see for example reference [123]), the isolation of FID emission from the echo transition is not possible. This is because transferring the input coherence to a different frequency and rephasing the ensemble requires a minimum of two rephasing pulses, and in a three-level system these two pulses inevitably address a common energy level. Therefore the coherence generated by the first pulse will be transferred to some extent by the second pulse onto the echo transition.

Aside from the immediate benefit of spectral distinguishability between FID noise and the echo emission, there are two other major benefits of this pulse sequence. Firstly, the 4LE involves two rephasing pulses, and as a result the phase-matching condition between input pulse, echo, and the π-pulses can be satisfied in geometries other than co-propagating (as will be shown in §6.3). The echo can then potentially be emitted in a different spatial mode to the bright pulses.

Secondly, the storage of light in atomic media is achieved by transferring coherence initially on an optical transition to the long-lived hyperfine ground states. This is an inherent feature of the 4LE sequence; provided the π-pulses are applied in the appropriate order, during the period τ_b (see figure 6.1) the coherence is stored between ground state levels.

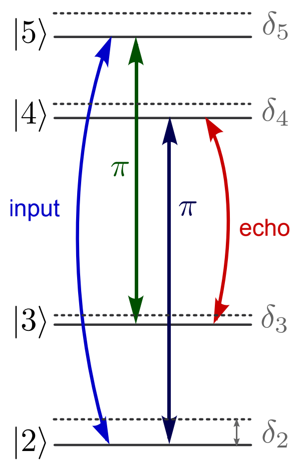

Figure 6.2: Energy levels addressed in the four-level echo sequence. The solid horizontal lines represent the energy levels resonant with the applied fields, whereas the dashed lines represent actual levels of an individual atom in the ensemble, with detunings $\delta_{2,3,4,5}$ relative to the applied fields.

6.3 Theory of rephasing using four levels

At first glance of the four-level echo pulse sequence, it is far from obvious why the atomic ensemble would rephase similarly to the two-level case. Unfortunately the elegant Bloch-sphere picture which provides insight into evolution of two-level systems does not retain its intuitive geometrical interpretation when applied to four-levels. To model the four-level echo process the evolution-operator formalism developed for the two-level case in §3.2 will be extended. Firstly, the matrix representation of the four-level basis states (see figure 6.2) is:

$$|\psi\rangle = \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix} = a\,|2\rangle + b\,|3\rangle + c\,|4\rangle + d\,|5\rangle \tag{6.1}$$

The system Hamiltonian is written as the sum of the atomic part, as well as four interaction terms corresponding to the four transitions marked in figure 6.2. It is assumed that the driving fields have a negligible effect on any transitions other than the one it is closest to. Transformation into the rotating-frame of all transitions simultaneously is performed using the unitary operator $\hat{U}(t) = \exp\left(-\frac{i}{\hbar}\mathscr{H}_{\text{resonant atom}}\right)$. After applying the rotating-wave approximation, and assuming the applied pulses are 'hard', then all the relevant evolution operators for each segment of the pulse sequence are found as;

$$\text{free evolution:} \qquad \hat{E}_\delta(t) = \begin{bmatrix} e^{-i\delta_2 t} & 0 & 0 & 0 \\ 0 & e^{-i\delta_3 t} & 0 & 0 \\ 0 & 0 & e^{-i\delta_4 t} & 0 \\ 0 & 0 & 0 & e^{-i\delta_5 t} \end{bmatrix} \tag{6.2}$$

$$\pi/2 \text{ pulse on } |2\rangle \rightarrow |5\rangle: \qquad \hat{E}_{25}^{\frac{\pi}{2}} = \begin{bmatrix} \frac{1}{\sqrt{2}} & 0 & 0 & -\frac{ie^{i\mathbf{k}_{25}\cdot\mathbf{r}}}{\sqrt{2}} \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -\frac{ie^{-i\mathbf{k}_{25}\cdot\mathbf{r}}}{\sqrt{2}} & 0 & 0 & \frac{1}{\sqrt{2}} \end{bmatrix} \tag{6.3}$$

$$\pi \text{ pulse on } |3\rangle \rightarrow |5\rangle: \qquad \hat{E}_{35}^{\pi} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -ie^{i\mathbf{k}_{35}\cdot\mathbf{r}} \\ 0 & 0 & 1 & 0 \\ 0 & -ie^{-i\mathbf{k}_{35}\cdot\mathbf{r}} & 0 & 0 \end{bmatrix} \tag{6.4}$$

$$\pi \text{ pulse on } |2\rangle \rightarrow |4\rangle: \qquad \hat{E}_{24}^{\pi} = \begin{bmatrix} 0 & 0 & -ie^{i\mathbf{k}_{24}\cdot\mathbf{r}} & 0 \\ 0 & 1 & 0 & 0 \\ -ie^{-i\mathbf{k}_{24}\cdot\mathbf{r}} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{6.5}$$

where for the operator $\hat{E}_T^A$, the resonant transition is denoted by T, and A is the pulse area (equal to the Rabi-frequency multiplied by the pulse length; $\Omega_T t$). For a single atom at position $\mathbf{r}$ initially in state $|2\rangle$, the final state after applying the entire pulse sequence is obtained as;

$$|\psi_{out}\rangle = \hat{E}_\delta(\tau_a) \times \hat{E}_{24}^{\pi} \times \hat{E}_\delta(\tau_b) \times \hat{E}_{35}^{\pi} \times \hat{E}_\delta(\tau_a) \times \hat{E}_{25}^{\frac{\pi}{2}} \times |2\rangle \tag{6.6}$$

$$= \frac{-1}{\sqrt{2}} \begin{bmatrix} 0 \\ e^{-i[\tau_a(\delta_3+\delta_5)+\delta_3\tau_b+(\mathbf{k}_{25}-\mathbf{k}_{35})\cdot\mathbf{r}]} \\ e^{-i[\tau_a(\delta_2+\delta_4)+\delta_2\tau_b+\mathbf{k}_{24}\cdot\mathbf{r}]} \\ 0 \end{bmatrix} \tag{6.7}$$

Finally, the density matrix corresponding to this state is determined[3] and compared to the state that would be obtained by applying a $\frac{-\pi}{2}$ pulse[4] (wavevector $\mathbf{k}_{echo}$) directly to the transition $|3\rangle \rightarrow |4\rangle$ for an ensemble initially in state $|3\rangle$. This represents the expected outcome of a two-pulse echo sequence on this transition alone. These density matrices are:

[3]For comparing two states, density matrices are used rather than state vectors simply because this removes the effect of any global phase shifts.

[4]A negative pulse area physically means the field has a 180° phase shift relative to the other pulses. In this model, which ignores evolution due to detuning while the fields are on, a pulse area of $\frac{-\pi}{2}$ is equivalent to a $\frac{3\pi}{2}$ pulse

$$\rho_{echo}(\mathbf{r}) = \frac{1}{2}\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & -ie^{i\mathbf{k}_{echo}\cdot\mathbf{r}} & 0 \\ 0 & ie^{-i\mathbf{k}_{echo}\cdot\mathbf{r}} & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{6.8}$$

$$\rho_{out}(\mathbf{r}) = \frac{1}{2}\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & -ie^{i[A\tau_a+B\tau_b+(\mathbf{k}_{24}-\mathbf{k}_{25}+\mathbf{k}_{35})\cdot\mathbf{r}]} & 0 \\ 0 & ie^{-i[A\tau_a+B\tau_b+(\mathbf{k}_{24}-\mathbf{k}_{25}+\mathbf{k}_{35})\cdot\mathbf{r}]} & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{6.9}$$

where $A = (\delta_2 + \delta_4) - (\delta_3 + \delta_5)$

$B = \delta_2 - \delta_3$

The coefficient A is the difference in detunings of the two π-pulse transitions. Atoms in the ensemble with $(\delta_2 + \delta_4) = (\delta_3 + \delta_5)$ will be re-phased. Therefore this echo sequence reverses the effects of inhomogeneity that is *common* to *both* optical transitions. With a non-zero delay time τ_b, there will be additional decoherence occurring due to inhomogeneity of the RF transition related to the term $e^{-iB\tau_b}$.

Optimal phasematching occurs when;

$$\mathbf{k}_{\text{echo}} + \mathbf{k}_{25} = \mathbf{k}_{24} + \mathbf{k}_{35} \tag{6.10}$$

This is clearly satisfied when all input pulses are co-propagating. However, unlike in the two-level two-pulse echo, there are alternate geometries that fulfill this phasematching condition. In particular, if the two π-pulses are counter-propagating ($\mathbf{k}_{24} = \mathbf{k}_{35}$) then the ensemble will radiate the echo in the opposite direction to the input pulse. (This counter-propagating regime will be used for demonstrating RASE in Chapter 7.)

In summary, according to this model the proposed four-level echo sequence does indeed rephase an initial coherence between two levels ($|2\rangle \leftrightarrow |5\rangle$) as an echo on a *different* transition ($|3\rangle \leftrightarrow |4\rangle$). However there is additional decoherence as compared to the two-level echo; during period τ_a due to any un-correlated inhomogeneity in the two π-pulse transitions, and during the period τ_b due to inhomogeneity on the RF transition. The latter of these can potentially be rephased, as will be seen in §6.5.2.

6.4 Four-level echo experimental demonstration

For this demonstration, a 2 cm long 0.005% $Pr^{3+} : Y_2SiO_5$ sample was used. A schematic of the experimental setup is shown in figure 6.3. To characterize the 4LE, we are only interested in a particular subset of ions that will be resonant with *all* transitions shown in figure 6.1(a), and would prefer not to excite ions in alternative frequency classes that are irrelevant for the experiment and would affect the propagation of the pulses. So in order to characterize the 4LE 'cleanly', we first tailor the spectrum to pump all but one frequency subgroup of ions into a dark ground state, where they will remain for the ground state cross-relaxation time of 100 s.

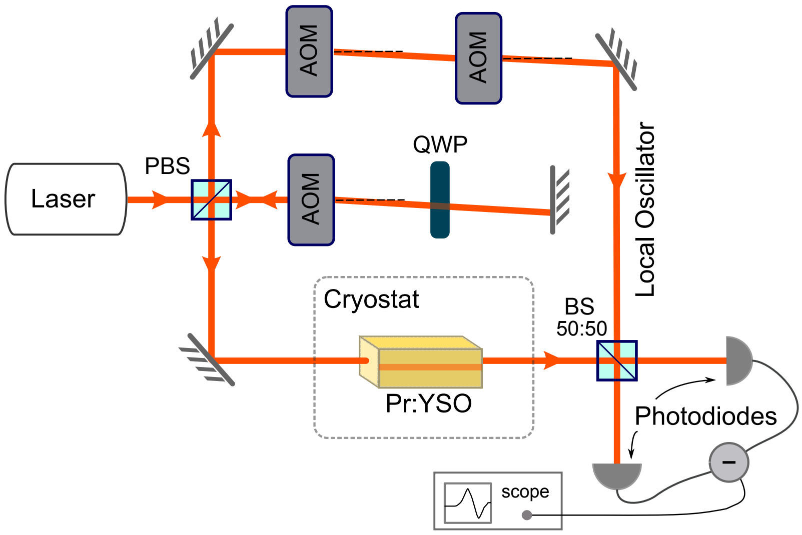

Figure 6.3: Experiment setup for testing the four-level echo. The 0.005% Pr^{3+}:Y_2SiO_5 sample is in a cryostat cooled to ∼2 K. The laser beam is incident on a polarizing beam splitter (PBS) and passes through two acousto-optic modulator (AOM) systems. One double-pass AOM generates all of the required frequencies for the feature preparation and the echo pulse sequences. The light that passes through the second set of AOMs in series is used as a local oscillator, which is mixed with the beam through the sample at a beam-splitter (BS) for balanced heterodyne detection. See Appendix F for detailed equipment specifications.

The spectral hole-burning process begins by iterating between 5 different frequency optical fields. Four of these fields are sweeping ±1 MHz centred at frequencies ω_{25}, ω_{35}, ω_{24} and ω_{34}, and the fifth field is at ω_{15}. The purpose of this initial step, similarly to techniques described in [88, 124], is to pump away any population that is not in the desired frequency subgroup for the experiment. Note that the bare minimum number of different frequency fields necessary for this step is 3; to drive transitions from each of the ground states. However it was seen here that the hole-burning was more efficient using all 5 frequencies; the four involved in the actual echo sequence, plus one additional field to address the remaining ground state.

Having selected out a particular frequency class of ions, the next step is to create a spectral feature in state $|2\rangle$. This is achieved by iterating between the ω_{35} and ω_{34} sweeping fields (to ensure state $|3\rangle$ remains empty) and applying the ω_{15} field to pump a narrow spectral population back into state $|2\rangle$. The power and duration of the ω_{15} field determines the width and optical depth of the feature. This field is single-tone, and the frequency-stabilized laser linewidth is < 1 kHz. However due to inhomogeneity between the RF levels the narrowest possible feature width burnt back will be on the order of 10 kHz. Typically the absorption feature in this experiment was 40 to 50 kHz wide, with the entire burning sequence taking ∼40 ms.

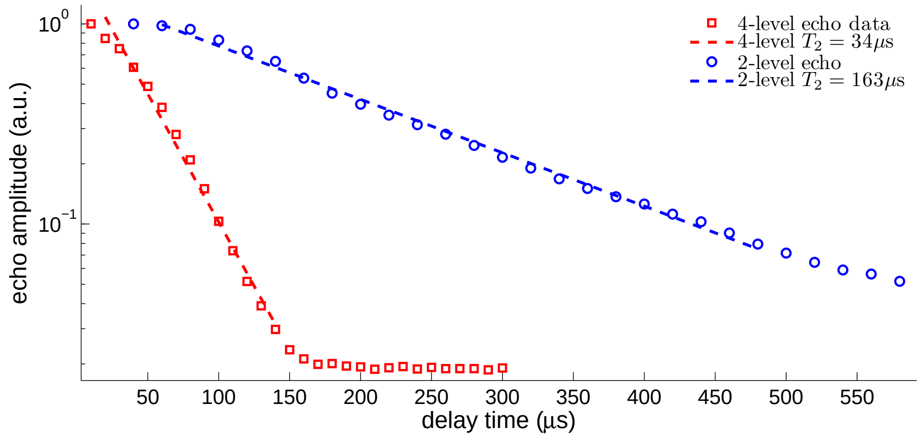

Figure 6.4: Measured photon echo amplitude with varying delay time. The x-axis represents the time between input pulse and the echo; $2\tau_a$ for the 4-level sequence (here $\tau_b = 0$) and 2τ for the 2-level sequence. Estimates of the delay times for both sequences are obtained by fitting an exponential to the data.

To ensure that the Pr energy levels are degenerate, a set of current-carrying coils was placed around the cryostat to null the residual magnetic field. Without this correction, there is a clear beat in the echo signal as delay time is varied, which suggests the ambient magnetic field induces a splitting of the hyperfine levels of approximately 40 kHz.

For the experimental results presented in this chapter, all beams were co-propagating.

6.5 Efficiency and coherence measurements

6.5.1 Optical coherence

The coherence time T_2 was determined for both sequences (using two or four levels) by measuring the echo amplitude with increasing delay time. Figure 6.4 shows the results for the optical transition. As discussed earlier (see §6.3), the two π-pulses used to generate the echo in the 4-level sequence will only rephase the component of the detunings that is common to both optical transitions (see equation 6.9). Due to this additional inhomogeneity, the coherence time is much shorter. The decay times were measured as $T_2^{4\mathrm{LE}} = 34$ µs and $T_2^{2\mathrm{LE}} = 163$ µs for the four- and two-level sequences respectively.[5]

6.5.2 Hyperfine coherence

One of the primary motivations for interfacing atoms and light is for long-term storage of light, which is ultimately achieved by transferring coherence initially on an optical transition to the long-lived ground states. This process is an inherent property of the 4LE sequence; the initial π-pulse transfers the coherence to the ground states, where it can be

[5] This coherence time for the 4LE was extended to ~ 80 µs for experiments described in the next chapter (see §7.5). This was due to improved cancellation of the residual magnetic field.

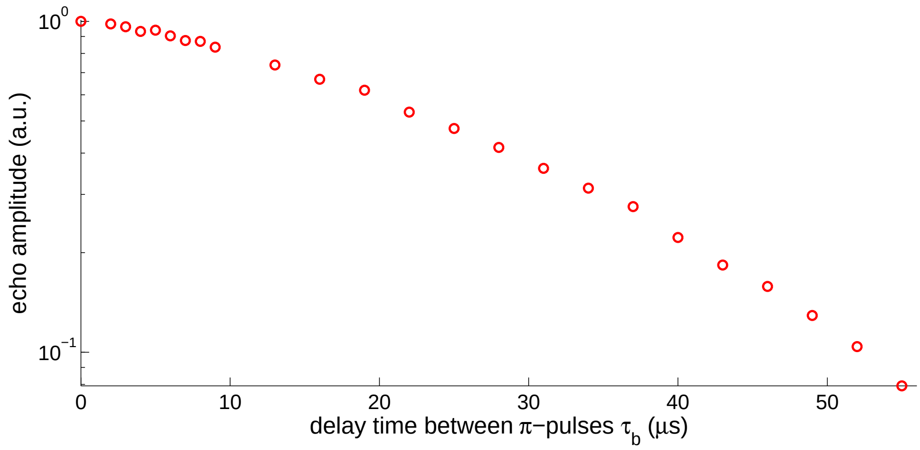

Figure 6.5: 4-level echo amplitude as a function of τ_b, the delay time between the two π-pulses. The loss of coherence is due to inhomogeneity on the $|2\rangle \leftrightarrow |3\rangle$ transition. The decay is not simply exponential, meaning the transition profile isn't purely lorentzian, however the inhomogeneity is estimated to be on the order of 10 kHz. For this data, $\tau_a = 15$ µs.

stored until the second π-pulse is applied to rephase on the optical transition.

The ground-state coherence decay in the 4LE sequence was determined by measuring the echo amplitude with increasing τ_b. Although the individual ions have long ground-state coherence times (0.5 ms in zero field), as seen in figure 6.5, the ensemble coherence decays much faster ($\sim$ 25 µs). This is due to inhomogeneous broadening on the $|2\rangle \leftrightarrow |3\rangle$ transition (see equation 6.9). The decay profile is not consistent with a pure lorentzian, nor a gaussian lineshape, however indicates a transition width the order of 10 kHz.

6.5.3 Extending the storage time

To take advantage of the long ground-state coherence times possible in Pr:YSO, it would be necessary to rephase this inhomogeneity on the RF transition. A coherence time on the order of 1 s is possible. This is achieved by applying a specific magnetic field to optimize the individual ion's coherence times. This is referred to as the 'critical-point' technique, discussed briefly in §2.4.3, and also in Chapter 9.

The pulse sequence proposed to rephase the RF inhomogeneity is depicted in figure 6.6. Note that only one RF π-pulse would be required to rephase the inhomogeneity between the ground states. However, for continued rephasing of the optical coherence after the final pulse is applied requires the applied RF to have a total area of an integer multiple of 2π. Applying a similar evolution-operator calculation as described by equation 6.7 for this entire pulse sequence, and neglecting the **k**-vector terms, the final density matrix describing the system at the time the echo is expected is;[6]

[6]Since the sample volume is $\ll \lambda^3$ for in the RF frequency range, these RF-pulses don't alter the phase-matching condition defined by the optical wave-vectors (equation 6.10).

$$\rho_{out} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & ie^{i[(\delta_2+\delta_4)-(\delta_3+\delta_5)]\tau_a} & 0 \\ 0 & -ie^{-i[(\delta_2+\delta_4)-(\delta_3+\delta_5)]\tau_a} & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{6.11}$$

All terms proportional to $(\delta_2 - \delta_3)$ and τ_c have been cancelled out, with only the optical detuning terms remaining. The only additional change is in the phase of the final echo. The coherence between states $|3\rangle$ and $|4\rangle$ is here similar to having directly applied a $\frac{\pi}{2}$ pulse at ω_{34} to an ensemble initially in state $|3\rangle$, that is the Bloch-vector will rephase on the opposite side of the Bloch sphere compared with the state described by equation 6.8.

Applying pulses directly to the $|2\rangle \leftrightarrow |3\rangle$ transition could potentially also transfer unwanted coherence between all transitions involved, and thus negate the original purpose of using the 4LE sequence (to remove FID emission from the echo mode). However the application of RF π-pulses is quite different from the optical equivalent, and there are two main reasons why the 4LE with RF rephasing can remain FID-free:

- It is technically much easier to apply spatially-uniform pulses in the RF frequency range compared to optical. Improved field homogeneity, and therefore Rabi-frequency uniformity across the ensemble, means that there is considerably less FID-emission generated on this transition.

- The *inhomogeneous* linewidth of the optical and RF transitions are similar, ~ 10 kHz, however the *homogeneous* linewidth is much smaller for the RF transition (1 Hz) compared to the optical (5 kHz). That is, the ratio T_2^*/T_2 for the RF transition is even larger than for the optical transition. This means that the delay time (τ_c in figure 6.6) can be very long compared to T_2^*. Therefore the FID-emission (which decays exponentially) will be negligible by the time the echo is rephased.

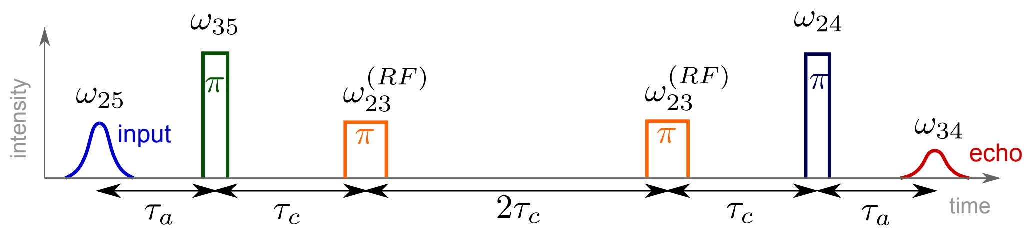

Figure 6.6: Proposed pulse sequence for 4LE including rephasing on the RF transition.

6.5.4 Echo efficiency

The efficiency of the 4LE relative to the 2LE was determined by preparing a feature of 54%±1% absorption (in terms of field amplitude), then measuring the echo for a weak input beam. An example measurement (in the frequency domain) is shown in figure 6.7.

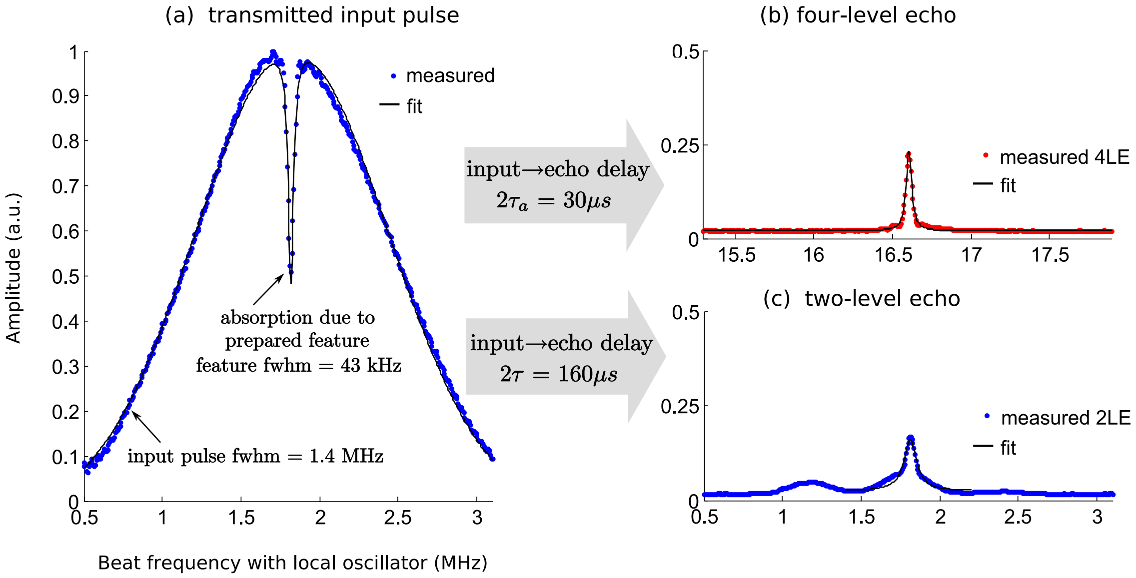

Figure 6.7: Efficiency measurements comparing the 4LE and 2LE sequences. (a) The transmitted input pulse; peak power is 5 µW, or 0.3% of the π-pulse peak power. The x-axis is the beat frequency of the signal mixed with a local oscillator beam. The local oscillator is centred 1.8 MHz from the input pulse. (b) The echo obtained using the 4LE sequence with a total delay time ($2\tau_a$) of 30 µs, and $\tau_b = 0$. The beat frequency between the local oscillator and the echo is 16.6 MHz. (c) The echo obtained using the 2LE sequence with a total delay time (2τ) of 160 µs. This echo occurs at the same frequency as the input pulse. All traces are the average of 400 shots.

The input-pulse peak power was 0.3% of the π-pulse peak intensity. Also, the input-pulse was spectrally much broader than the prepared feature such that most of the pulse was transmitted. Fitting the envelope of the transmitted pulse (as a narrow lorentzian subtracted from a broad gaussian peak) gives the input amplitude as well as the optical depth and width of the prepared spectral feature. The delay times for each sequence were chosen to ensure the echo could be temporally resolved from the FID following the π-pulse(s). However there is significant background in the echo spectrum for the two-level case (figure 6.7(c)). This is attributed to FID emission, and also some echo emission from background population in the region surrounding the feature.

The efficiency is determined as the amplitude of the echo normalized to the amplitude of the input pulse, and the results are summarized in table 6.1. Since the echo is measured at a different delay time for each sequence, the resulting efficiency is extrapolated back to zero delay time using the previously measured coherence times. The result shows that the 4LE is 5% less efficient than the 2LE. This slight difference can be accounted for by considering the propagation of the echo. The propagation of the input pulse is identical for the two pulse sequences, however the rephasing occurs on different transitions, of different oscillator strengths. The relative oscillator-strengths[7] of $|2\rangle \leftrightarrow |5\rangle$ and the $|3\rangle \leftrightarrow |4\rangle$ transitions are 0.60 and 0.55 respectively [88]. This translates to a relative difference in transition dipole moment[8] of 4.3%, and thus the echo emission in the four-level scheme

[7] related to field intensity
[8] related to field amplitude

	2-level echo	*4-level echo*
input pulse absorption ±1% (at feature frequency)	54%	54%
echo efficiency ±5% (echo amplitude/input pulse amplitude)	13%	21%
delay time of efficiency measurement	$2\tau = 160\mu$s	$2\tau_a = 30\mu$s $\tau_b = 0\mu$s
scale factor (accounting for decoherence during delay time)	0.21	0.37
scaled echo efficiency ±5% (at zero delay time)	**63%**	**57%**

Table 6.1: Summary of efficiency measurement results comparing 4LE and 2LE

experiences 4.3% less gain as it propagates relative to the 2LE. This is the order of the measured discrepancy in efficiencies. Therefore we conclude that the four-level echo has similar efficiency to the two-level equivalent. This is to be expected given the dynamics of the rephasing are essentially the same as discussed above in section 6.3.

6.5.5 Phase of the echo

Heterodyne detection is a phase-sensitive technique, and can be applied here to check the phase response of the 4LE to changes in the phases of the input or rephasing pulses. The model developed in §6.3 can be applied to predict the phase-related behaviour of the 4LE. To include a phase variable, the description of the classical applied field (equation 3.12) is re-written as $\mathbf{E}_f = \mathcal{E}\cos[\nu_f t - (\mathbf{k}_f.\mathbf{r} + \phi_f)]$. Here f labels the different frequency fields, and the factor ϕ has been added. This phase is included in the density matrix of the final state (equation 6.9) to add independent phase factors for all applied pulses and the echo by making the following transformations:

$$\begin{aligned}
\text{input pulse:} &\quad \mathbf{k}_{25}.\mathbf{r} \to \mathbf{k}_{25}.\mathbf{r} + \phi_{in} \\
\text{first } \pi\text{-pulse:} &\quad \mathbf{k}_{35}.\mathbf{r} \to \mathbf{k}_{35}.\mathbf{r} + \phi_{\pi 1} \\
\text{second } \pi\text{-pulse:} &\quad \mathbf{k}_{24}.\mathbf{r} \to \mathbf{k}_{24}.\mathbf{r} + \phi_{\pi 2} \\
\text{echo:} &\quad \mathbf{k}_{34}.\mathbf{r} \to \mathbf{k}_{34}.\mathbf{r} + \phi_{echo}
\end{aligned}$$

Then equating the modified phase factors in equations 6.8 and 6.9 gives:

$$\phi_{echo} = \phi_{\pi 1} + \phi_{\pi 2} - \phi_{in} \tag{6.12}$$

This means that a relative change in phase of $\Delta\theta$ in either π-pulse will change the phase of the echo by the same amount $\Delta\theta$, whereas changing the phase of the input pulse alters the phase of the echo in the opposite direction, $-\Delta\theta$. The equivalent condition for the two-level echo is similarly obtained as;[9]

$$\phi_{echo} = 2\phi_{\pi} - \phi_{in} \tag{6.13}$$

Figure 6.8 shows the measured phase of the 4LE for four different phases of the input pulse, while the phases of the π-pulses remained fixed. For all results presented in this section, the spectral trenches burned during the feature preparation stage were increased from 1 to 2MHz in width to accommodate the addition of reference pulses. The process of measuring the phase of the heterodyne detection system using reference pulses is described in Appendix C.

In figure 6.8, the blue lines are the histogram of the measured echo phase. These polar plots can equivalently be interpreted as the transverse ($\hat{x} - \hat{y}$) plane of the Bloch sphere (see figure 3.1) in the rotating frame of the two-level subset ($|3\rangle$ and $|4\rangle$) of the four energy levels. The red line represents an overlay of the Bloch vector for the input pulse in the rotating frame of the input-pulse frequency. The relative phase between these two rotating-frames can be set arbitrarily, and has been chosen to highlight the similarities in behaviour of the four- and two-level echo sequences. As predicted by equation 6.12, the measurement results shown in figure 6.8 show the echo phase rotates in the opposite direction to the input pulse phase.

The noise in the phase measurement was typically 2.5°. There was a small phase-shift that accumulates across the sequence of shots, resulting in a net offset of 5.5° in the 6 minute period of time recording all the data shown in figure 6.8. This is attributed to a slow drift in the clock frequency of the digital oscilloscope used for data acquisition relative to the RF generators driving the AOMs for the optical pulses.

Figure 6.9 shows the response of the echo phase to changes in the phase of the first π pulse. Aside from the same 5.5° shift seen in figure 6.8, the Bloch-vector representing the echo rotates in an identical fashion to the π_1 pulse. This is consistent with what is predicted by equation 6.12.

Aside from the phase of the input and rephasing pulses, a major factor affecting the phase of the echo are the propagation effects. All of the input pulses as well as the echo itself will experience dispersion as they propagate through the ensemble of resonant ions. To demonstrate this effect, the area of the rephasing pulses were varied from 0.43π to π, and as shown in figure 6.10, the net shift in the phase of the resulting echo is 80°. However the sensitivity of the echo-phase to changes in the initial feature absorption were seen to be minimal; as the feature absorption was varied from 10% to 90%, there was less than 6° variation in phase of the echo. Due to the plethora of parameters that will affect the propagation of light through the sample, the only meaningful relative phase measurements

[9] Although the phase behaviour of the 2LE was not measured here for the optical transition, it has been observed previously (for example see reference [124].)

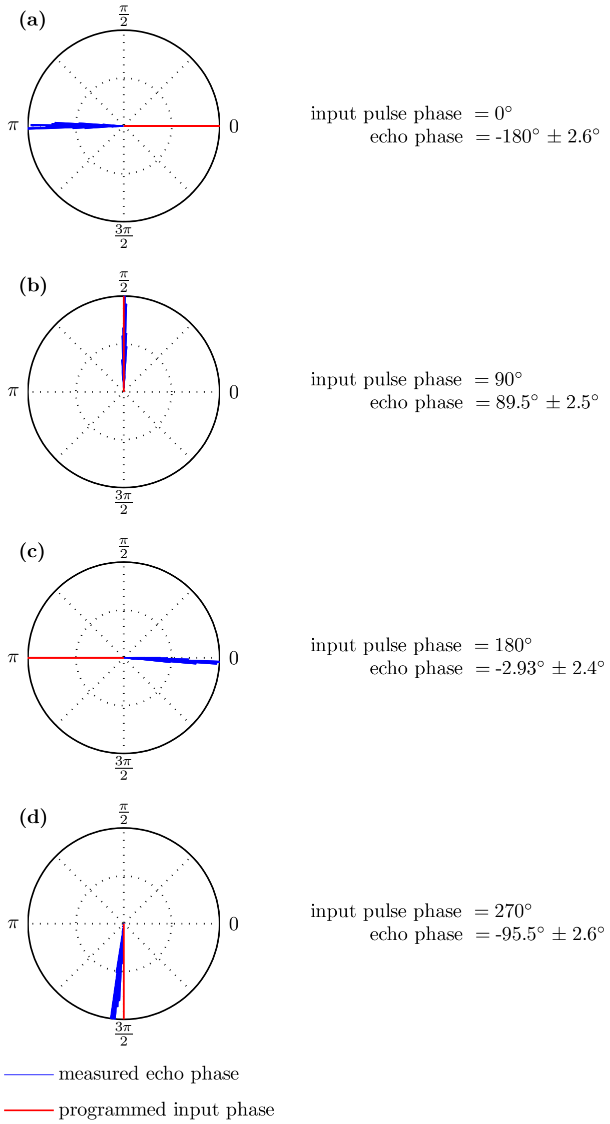

Figure 6.8: Four-level echo phase measurements. These are polar plots of the phase of the measured echo signal (blue), with varying phase of the input pulse (marked in red). The graphs can be interpreted as the transverse planes of two overlayed Bloch-spheres, one in the rotating frame of the echo-transition, and the other in the rotating frame of the input-pulse transition. The phase between the two rotating-frames can be set arbitrarily, and has been chosen such that the echo signal has a phase of π when the input pulse phase is zero. Similarly to the two-level case, the echo Bloch-vector rotates in the opposite sense to the input pulse. The radial dimensions of each plot have been individually normalized to the maximum of the echo-phase histogram, and for each phase setting there was ∼200 shots recorded. There are occasional glitches in the phase measurement (>5 standard deviations from the mean), and these data points have been omitted from the histograms (<0.4% of the total data points).

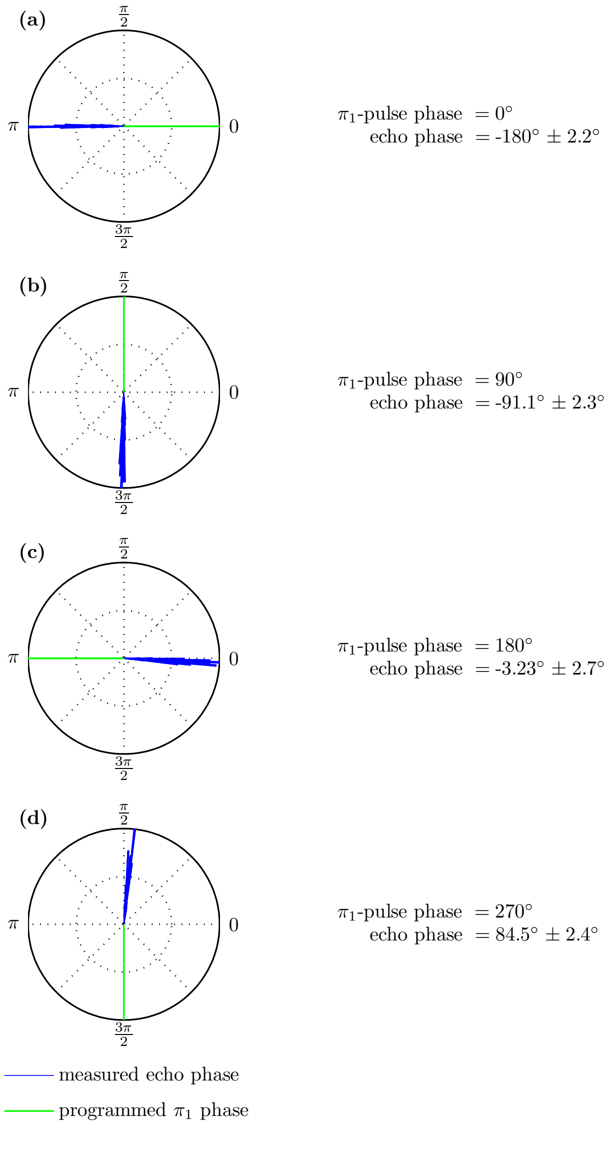

Figure 6.9: Four-level echo phase measurements. These are polar plots of the phase of the measured echo signal (blue), with varying phase of the first π-pulse (marked in green). The input pulse and second π-pulse phases are fixed. This shows that the echo phase rotates identically to the phase of π_1. Data acquisition and processing is identical to the description given for figure 6.8.

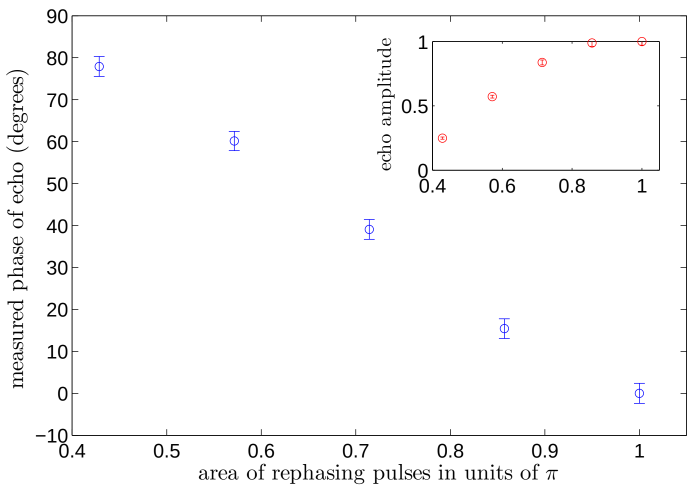

Figure 6.10: Measuring the sensitivity of the 4LE phase to different pulse-areas of the two rephasing pulses. The phase is changed due to the different propagation properties of the ensemble. Inset shows the amplitude of the echo, normalized to the optimum which occurs when the rephasing pulses have an area of π.

are between data sets where the absorptive properties of the ensemble, and all of the input pulse strengths, spectra, and temporal widths are kept constant.

6.6 Noise in the echo mode

The primary motivation for investigating this four-level echo sequence is to rephase atomic coherence at a frequency that is otherwise completely dark; a major technical challenge in the two-level case simply due to imprecise π-pulses and their associated FIDs. To verify that the 4LE is as quiet as expected, spectra were recorded for the 4LE sequence with no input pulse. The spectral intensity in the echo mode was then determined (the temporal mode is shown in figure 6.11). This spectral intensity was compared to the background noise level, which was determined using two different methods. The first approach was to use a different temporal mode. This alternative temporal envelope was generated simply by shifting the 'signal' temporal mode to a longer time after the π-pulses (starting at 45 μs instead of 0 μs in figure 6.11). The second method was to shift the spectral mode outside the 50 kHz bandwidth of the echo mode. Specifically, the spectral intensity was calculated at $\pm$100 kHz from the echo beat frequency of 16.6 MHz. Both of these methodologies give the same value of the background noise level within the uncertainty of the measurement.[10]

This background spectral intensity is comprised of two components; the shot noise in the optical field, and electrical noise in the detector. For the detector used in this experiment, 75% of the background noise was the shot noise (with 0.6 mW power in the local oscillator beam and at 16.6 MHz).

[10] A much more straightforward means of determining the background noise level would have been to simply use the period *before* the rephasing pulses were applied. Unfortunately this section of the signal was not recorded due to ignorance of the experimenter.

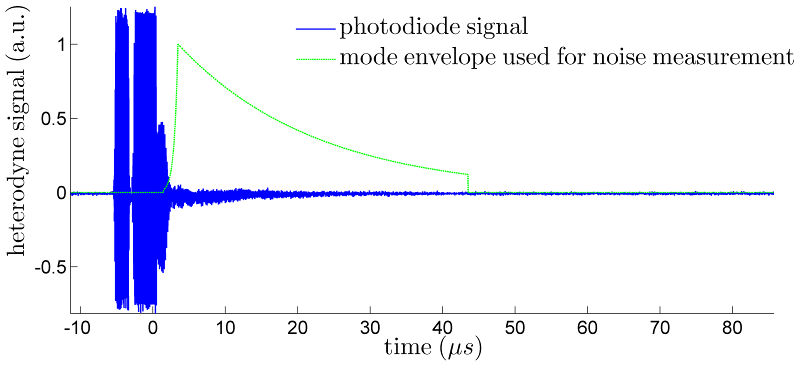

Figure 6.11: The temporal envelope used to calculate the noise in the echo mode (green), overlayed on a single recorded trace from the photodiode (blue). This mode is an exponential decay with a time constant of 19 μs, corresponding to the decay of the echo intensity. The two rephasing π-pulses occurring just before 0μs saturate the detection system. There is clearly noise evident in the trace following these π-pulses, however this is almost entirely at the frequencies of the π-pulses themselves.

Averaging across 7500 shots, the noise in the echo mode was found to be 1.04 ± 0.02 of the shot noise level.

Perhaps a more meaningful measurement would be relative to a number of photons, rather than the shot noise level. Although we are only interested in one field mode, in this heterodyne detection system there are multiple modes contributing to shot noise in this bandwidth. Firstly, frequencies detuned from the local oscillator by $\pm\Delta$, will *both* contribute to the measured spectrum at beat frequency Δ. Also, in this experiment there was a mismatch between detector and the digital data-acquisition-system bandwidths giving rise to aliasing. This effect was quantified, and it was found that 50% of the background noise (at 16.6 MHz) is in the lowest order beat frequency, and the remaining 50% is due to higher orders being aliased. Therefore, in this detection system, 25% of the shot noise corresponds to one photon in the echo mode. Using this conversion, it is estimated that there are approximately 0.2$\pm$0.1 photons on average per shot added to the echo mode.

The source of this added noise is amplified spontaneous emission; the second π-pulse at ω_{24} creates a population inversion, and some of these atoms will spontaneously relax to state $|3\rangle$, emitting light at the echo frequency ω_{34}. Considering the temporal envelope used, and the relative branching ratio, the amplification factor of this emission (i.e. what would be measured if this gain feature had been probed) is $\alpha L \approx 1.2$. The initial feature is known to be 54% absorptive, and the corresponding optical depth is $\alpha L \approx 1.6$. Thus approximately 75% of the ions were inverted, which is consistent with the expected quality of the π-pulse.

A calculation of the noise added in the same 20 µs temporal mode, but at the frequencies of the π-pulses gives a measure of the FID emission. The number of photons added in the spectral modes of the first (ω_{35}) and second (ω_{24}) π-pulses are 1600 and

80000 respectively. The magnitudes are disparate because the first pulse interacts with relatively few ions compared to the second.

6.7 Summary

The conclusion drawn from these measurements is that the level of noise added in the echo mode in the 4LE sequence, 0.2 photons/shot, is consistent with the level of amplified spontaneous emission expected. With frequency-resolving detection the FID emission in the same temporal mode, on the order of $\sim 10^5$ photons/shot has been avoided. In a two-level photon echo sequence, the best sensitivity that would be *theoretically* possible (i.e. with a perfect π-pulse) is at the level of the spontaneous emission. By rephasing in a four-level system, despite having imprecise π-pulses we are able to actually measure at this level of sensitivity.

A model of this four-level system was developed, and applied to determine the detuning components responsible for the reduced coherence time. In addition, the predicted phase relationships between the echo and the input/rephasing pulses were experimentally verified.

CHAPTER 7

RASE experiment

In this chapter we apply the four-level echo (4LE) sequence (introduced in chapter 6) to rephase amplified spontaneous emission. The aim of this experiment is to confirm the theoretical prediction that spontaneous emission can indeed be rephased, and to investigate the use of the photons produced as a source of non-classically correlated photon pairs. Ultimately photon-counting detection is used to investigate the correlations, however heterodyne techniques are also employed for characterizing the system parameters.

This experiment is fundamentally a search for intensity correlations between individual spatio-temporal-spectral field modes. As such, the dimensionality of the experimental setup is complicated, spanning the spatial, temporal and frequency domains.

Initially, in §7.1, the proposed experiment is briefly outlined. This mainly focuses on the differences between the 4LE experiment (with heterodyne detection) described in the last chapter, and the RASE experiment (with photon-counting detection) that will be presented in this chapter. Following this synopsis, the remainder of the chapter will be outlined.

7.1 Synopsis

In the RASE protocol, the ensemble begins in the excited state, and coherence is generated when a spontaneous emission event occurs. When this spontaneously emitted photon is detected, it heralds an entangled state of the ensemble since we have no knowledge of exactly which atom emitted the photon. Due to inhomogeneity across the ensemble, this coherence will dephase, and can be rephased similarly to a more conventional coherent superposition state, using photon echo techniques.

7.1.1 RASE with the four-level echo

The four-level echo was introduced in the previous chapter (Chapter 6). An input pulse generated a coherent state of the atomic ensemble at a particular frequency, and using two π-pulses this coherence was rephased on a different transition. The pulse sequence is modified slightly for the present experiment. Initially a π-pulse is applied to invert the ensemble, and a spontaneous emission event generates the ensemble coherence rather than an input pulse. A simplified schematic of the four-level echo pulse sequence applied to rephase spontaneous emission is shown in figure 7.1.

The main advantage of using the four-level echo sequence is that the coherent emission due to FIDs associated with imperfect π-pulses are at *different* frequencies to both the ASE and RASE fields. Also the phase-matching condition relating the ASE, RASE, and the rephasing π-pulses can be satisfied in a beam geometry other than co-propagating (see equation 6.10). This means that the FID emission can also be spatially separated from the ASE and RASE modes.

As an added bonus, in the interim period between the two rephasing π-pulses in the 4LE sequence, the coherence is stored on a ground state spin transition, which can have very long coherence times. The ground state coherence times can be extended to seconds using the critical-point technique; 4 orders of magnitude longer than the $\sim$100 μs long coherence time of the optical transition. Storing coherence between the hyperfine ground states will be necessary for any future applications which rely on probabilistic processes, including for instance the entanglement-swapping procedure involved in making a quantum repeater (see figure 1.2).

In many quantum-information-processing protocols in rare-earth ion ensembles, this shelving of the atomic coherence in the hyperfine ground states is an auxiliary process, whereas here it is an inherent property of the four-level echo.

7.1.2 RASE with photon-counting detection

The results from Chapter 6 verified that the four-level echo in a co-propagating geometry, allows for the detection of photon echoes down to the spontaneous emission sensitivity level using heterodyne detection methods. Heterodyne is an extremely useful diagnostic tool with high dynamic range, phase-sensitivity, and most importantly, spectral information. However, our long-term aim for the RASE protocol is to produce on-demand single photons, or alternatively generate heralded entangled atomic states. Direct measurement of the single-photons is preferential over heterodyne for this purpose. Therefore avalanche photodiode (APD) detectors are used for the bulk of the results presented in this chapter, with the heterodyne system used for some initial diagnostic measurements.

Using single photon detection in the RASE experiment poses some technical challenges as compared to heterodyne, largely because photon-counting provides no frequency discrimination. As a result there are two major modifications required to enable signal photons to be resolved from other emission.

The first, and most important modification, is to use a counter-propagating beam geometry in the four-level rephasing. The benefit of moving to an off-axis geometry is that the FID emission is *spatially separated* from the ASE and RASE modes. The co-propagating echo measurements of the last chapter found that the coherent emission associated with the FID was $\sim$6 orders of magnitude larger than the amplified spontaneous emission level.

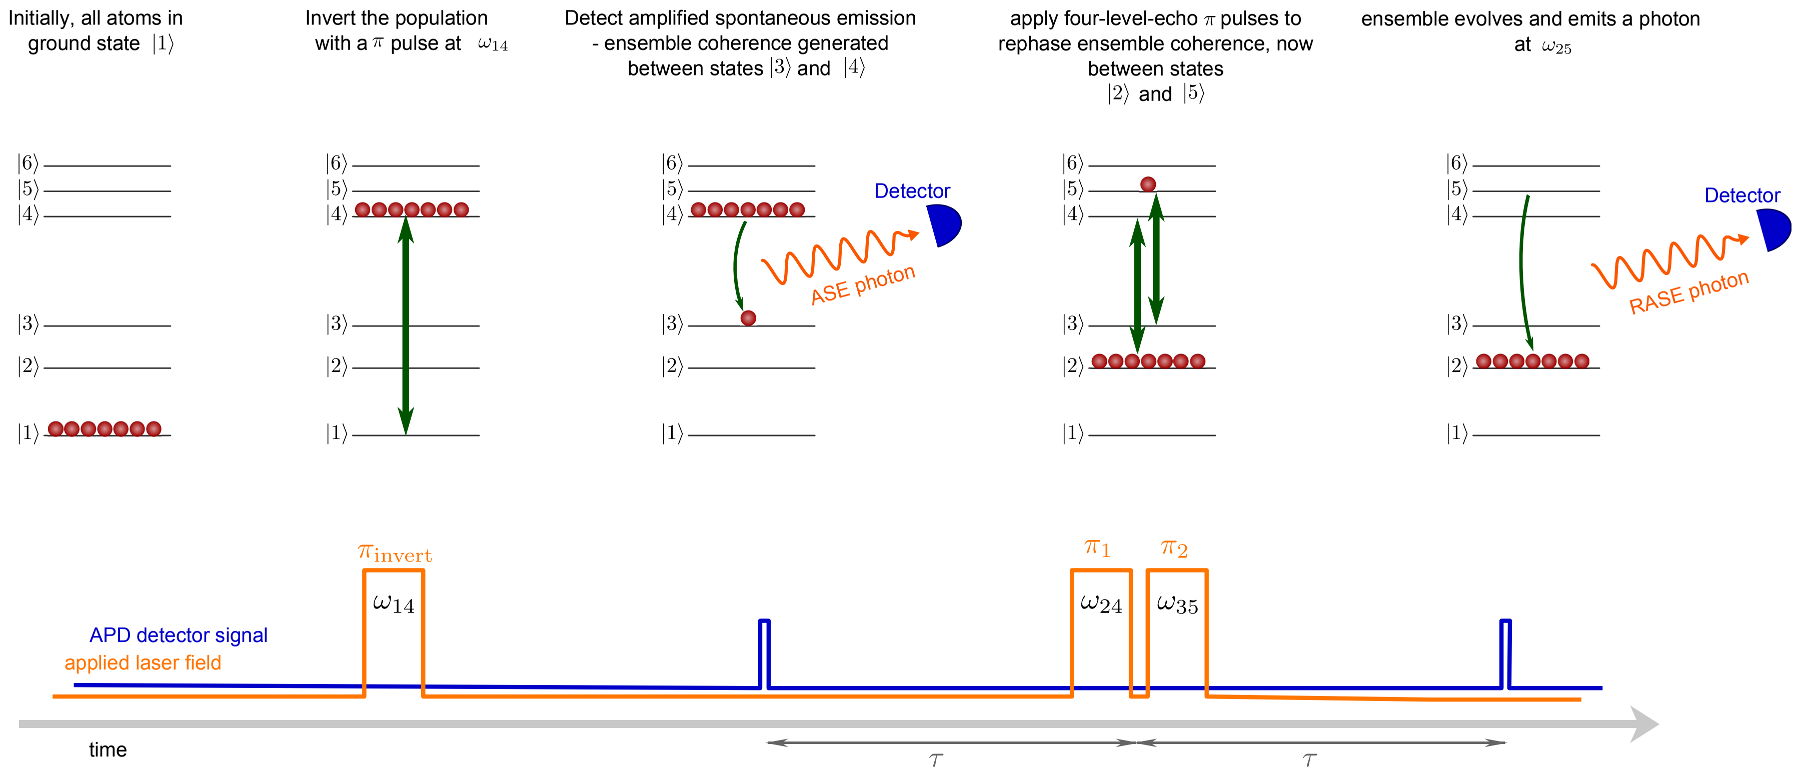

Figure 7.1: The pulse sequence used to generate amplified spontaneous emission, and rephase this spontaneous emission on a different transition. This is performed with the four-level echo pulse sequence introduced in chapter 6.

Thus, at the cost of significant experimental overhead, the counter-propagating beam geometry gains many orders of magnitude in noise reduction.

The second major modification in moving to photon-counting detection is to implement auxiliary frequency filtering. This is achieved using hole-burning techniques in a second region of the sample, where the spectrum is tailored to absorb any emission which isn't at the signal frequency. Given that the coherent FID emission is already removed from the detector spatial modes, the remaining noise emission in the signal modes is due to spontaneous emission between different hyperfine levels in the excited and ground state manifolds. Thus this noise emission is comparable in intensity to the ASE signal.

The implementation of this auxiliary spectral filtering requires significant experimental resources, and therefore only the RASE mode is filtered. The addition of any noise to the RASE mode will be more detrimental to the correlation fidelity than the same amount of noise added in the ASE mode, simply because the efficiency of the rephasing process is <100%.

7.1.3 Adjusting the ASE spectra via the Stark-shift

These two alterations; changing the beam geometry to be counter-propagating, and introducing auxiliary frequency filtering; are vital modifications of the experiment in implementing a photon-counting detection system. The resulting geometry involved 5 beam paths intersecting in various regions of a single 20 mm long crystal.

A final enabling factor, which arises as a consequence of the chosen sample and beam geometry, was to use a field-gradient to engineer the ensemble inversion (via the Stark shift). This was necessary so as to avoid significant stimulated emission (superradiance) from rapidly depleting the gain.

7.1.4 Chapter outline

The next section in this chapter, §7.2 will describe the experimental setup in detail. This includes the beam geometry and phase-matching considerations, details of the optical layout, the pulse sequence, the method of generating the initial inversion, the Stark-shift properties, and the sample holder details.

The next sections describe measurements of various experiment parameters that were necessary before the RASE protocol itself was attempted. In section 7.3, measurements quantifying the spectral properties of the inversion feature, and the rate of population decay, are described. Section 7.4 outlines the requirements for the frequency filtering, and presents measurements of the filter performance. Section 7.5 describes the echo efficiency in the counter-propagating geometry.

The latter sections present an analysis of the RASE experiment results with single photon counting detection. Firstly in §7.6 the count histograms are used to determine the signal and noise components of the detected intensity distributions. In section 7.7 the photon arrival times measured in each shot are used to calculate the second-order correlation function, and it is verified that amplified spontaneous emission is being successfully rephased. As shown in 7.8, although the correlation does not confirm the entangled nature of the photon pair, the degree of correlation measured is fully understood in terms of the experiment signal and noise levels and loss parameters. Finally in §7.10 the changes required that would enable a measurement of a non-classical correlation are outlined.

7.2 Experiment setup

7.2.1 Beam geometry

The spatial layout of the four-level echo demonstration in Chapter 6 involved the input pulse and both rephasing pulses co-propagating. However as shown by equation 6.10, it is possible to have perfect phase-matching without all beams propagating in the same spatial mode. In the context of the RASE experiment, the phase-matching condition is;

$$\mathbf{k}_{\text{RASE}} + \mathbf{k}_{\text{ASE}} = \mathbf{k}_{\pi_1} + \mathbf{k}_{\pi_2} \tag{7.1}$$

For this demonstration, a beam geometry where $\mathbf{k}_{\pi_1} = -\mathbf{k}_{\pi_2}$ and $\mathbf{k}_{\text{RASE}} = -\mathbf{k}_{\text{ASE}}$ (with $\mathbf{k}_{\text{RASE}|\text{ASE}} \neq \mathbf{k}_{\pi_1|\pi_2}$) was chosen. This *counter-propagating* beam geometry represents the simplest arrangement to align in which the bright rephasing pulses are spatially separated from the detection modes.

The beam geometry is depicted in figure 7.2 along with a schematic of the addressed transitions. The two rephasing π-pulses propagate in opposite directions along the path labelled *c*-*d*, and the avalanche photodiode (APD) detectors, labelled as D_A and D_R respectively, count the ASE and RASE photon-events along path *a*-*b*. For the entire optics schematic, as well as equipment details and specifications, see Appendix F.

7.2.2 Optical layout

The angle between the ASE/RASE detection paths (path *a*-*b*) and the rephasing-pulse spatial modes (path *c*-*d*) is approximately 10°. The ASE/RASE modes have a radius of 95 μm at the focus; measured with a beam-profiler. The π-pulse mode radius is 155 μm. Therefore the overlap region between these two modes in the sample is calculated to be ~2.4 mm long.

It is necessary to spectrally hole-burn along path *a*-*b* to ensure transmission of ASE and RASE photons through the entire length of the sample, which is why there is a beam-splitter in the ASE detector path; to allow for a mode-matched beam input along path *b*. Also, when the beam along path *b* is switched on to hole-burn this transmission window, the mechanical shutter (see §F.1) in front of the RASE detector is closed to prevent the excitation of long-lived fluorescence thought to originate from the glue in the optical fibre connector.

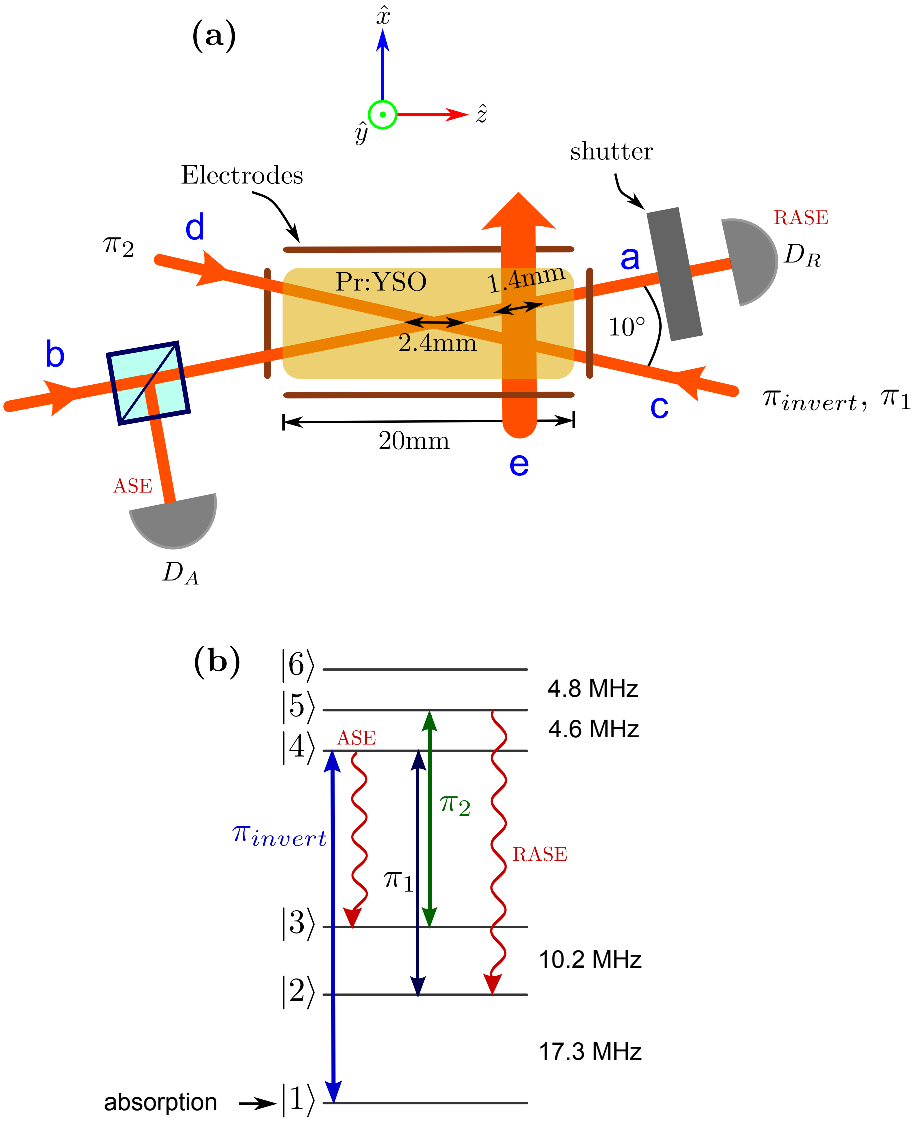

Figure 7.2: Overview of key experiment parameters for demonstrating RASE with the 4LE. (a) Counter-propagating beam geometry, with path labels a through e defined. The ASE and RASE modes are detected using APDs labelled D_A and D_R respectively. There is a mechanical shutter before the RASE detector, and a 50:50 beam-splitter before the ASE detector to allow for spectral hole-burning along path a. (b) Transitions used for the RASE experiment. The key difference between this and the 4LE experiment is that the sequence begins with a π-pulse at ω_{14} to populate excited state $|4\rangle$, and then the spontaneous emission from this level generates the ensemble coherence to be rephased.

The final beam-path (path e) is not associated with the four-level echo; it is used to implement the frequency filtering. This beam is 1.4 mm in diameter, and is used to tailor the spectrum in the overlap region between paths 5 and 1-2. This filter is designed to absorptively attenuate particular frequency modes that would otherwise contribute noise to the RASE signal. The specifics of the filter-burning process will be described in more detail below (§7.4).

Additional frequency-selective components in the optical setup are 10 nm bandpass interference filters located before each of the APDs to transmit the emission at 606 nm, and block the emission to intermediate crystal field levels at longer wavelengths (filter specifications given in §F.1).

Beam paths c, d and e propagate entirely in free-space, while the modes along paths a and b are single-mode-fibre coupled for interfacing with the detectors. In addition to reducing the background noise in detection, this fibre coupling also allows for the mode-matching to be easily quantified.

The optical frequencies along paths b, c and d were controlled using either two AOMs in series, or a double-pass AOM configuration. The frequency of the light along path e was controlled using an EOM in combination with a double-pass AOM to allow for more flexible frequency control.[1] The full optics schematic is shown in §F.1.

7.2.3 Detection efficiencies

The detection probability along the path from the interaction region in the sample to the RASE detector (path a) is 10.5%. This is determined by measuring the fraction of power transmitted from before the cryostat through to the detection fibre output, and multiplying this by the quantum efficiency of the APDs (67%).

Aside from the detector efficiency, the remaining losses are attributed to the cryostat windows (70% transmission), the bandpass interference filter (55% transmission), fibre coupling (60% efficiency) and reflection losses associated with propagation through a handful of lenses and mirrors (see figure F.1). The detection efficiency of the ASE mode is calculated to be 5.8%, experiencing similar transmission losses to the RASE mode, but with the inclusion of a 50:50 beam-splitter.

7.2.4 Ensemble preparation and pulse sequence

Figure 7.3 outlines the different stages of the experiment sequence, with each subfigure showing the *spatial* layout of the different *frequency* optical fields applied at a particular point in the *temporal* sequence. The process begins by hole-burning spectral-trenches 2 MHz wide centred on the frequencies of the π-pulses along path d, and similarly along path b burning 2 MHz wide trenches for transmission of the ASE and RASE frequencies. In the overlap region this ensures that all population in states $|2\rangle$ and $|3\rangle$ in the desired frequency-class are pumped into state $|1\rangle$.

[1] Also the laboratory supply of RF sources for driving the AOMs had at this point been fully exhausted.

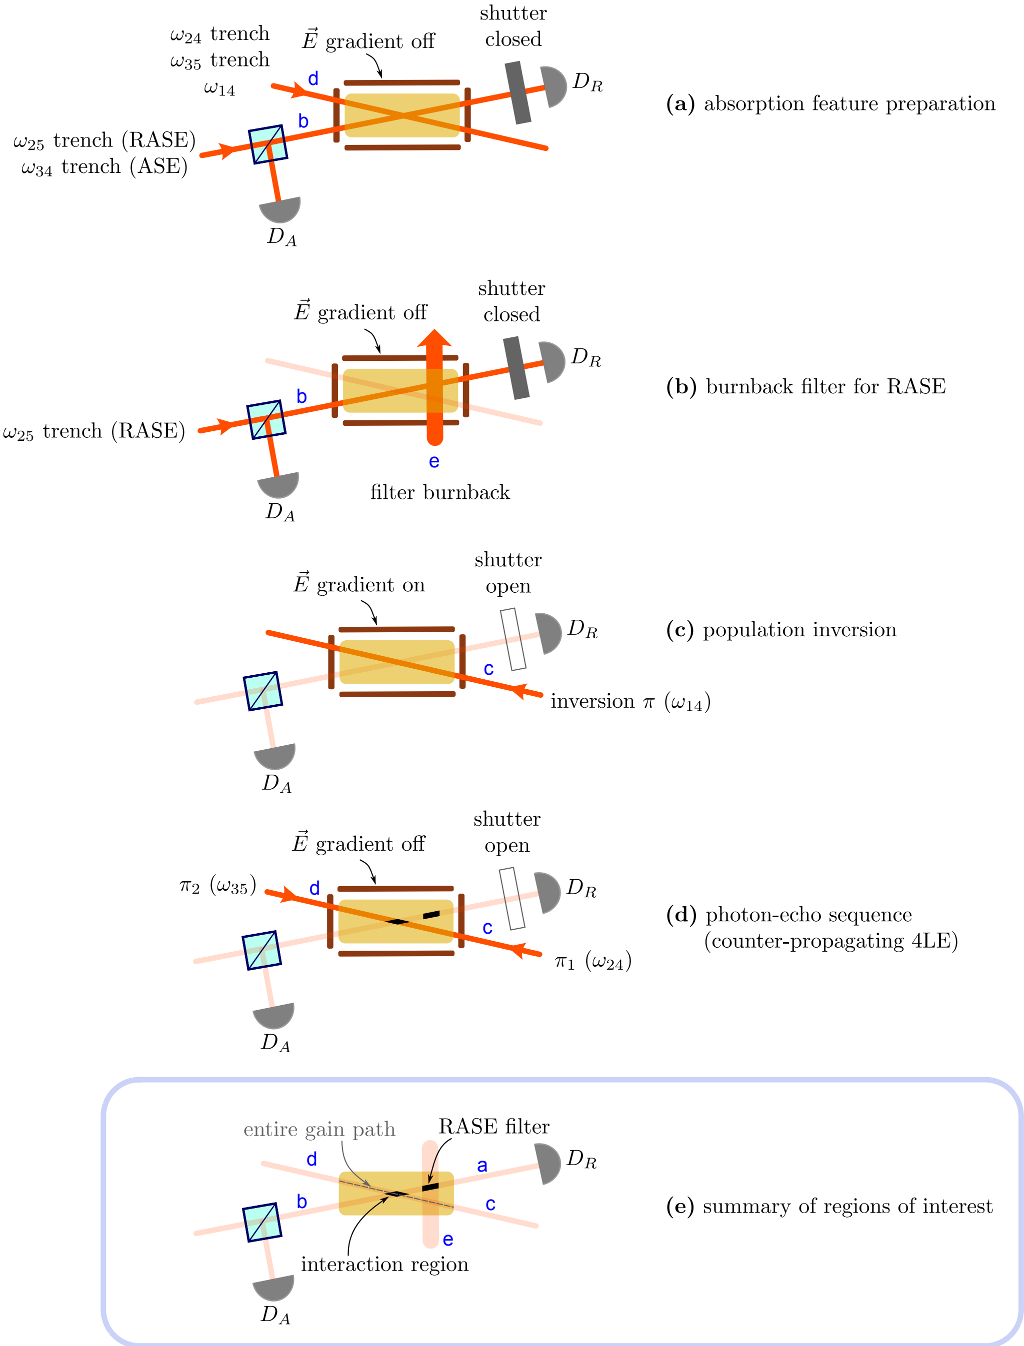

Figure 7.3: Overview of a single-shot of the RASE experiment. Each subfigure represents a different stage in the temporal sequence. Parts (a) and (b) depict the applied field configurations for the ensemble and frequency-filter preparation. (c) shows how the population inversion is generated, and (d) the counter-propagating four-level echo sequence. (e) is a summary of the important spatial regions in the sample. D_A and D_R are APD detectors for the ASE and RASE modes respectively, and the beam-splitter before D_A is 50:50.

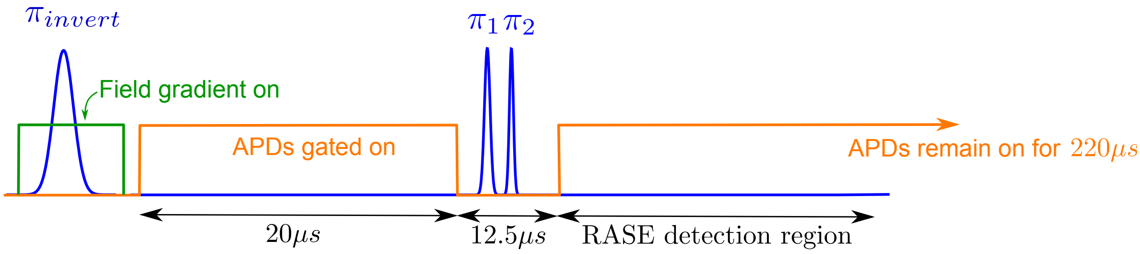

Figure 7.4: Diagram showing temporal sequence of the RASE experiment. The field-gradient is switched on during the inversion π-pulse. The APDs are gated off during all bright pulses and are activated for 20 μs following the inversion pulse, and 220 μs following the two rephasing π-pulses.

During the second ensemble-preparation step, shown in figure 7.3(b), the scanning field centered at the RASE frequency remains on, while the field along path e switches between four frequencies to transfer atomic population to absorb at three specific frequencies that would otherwise contribute noise to the RASE mode. The properties of this filtering process will be quantified in §7.4. The time taken for the entire ensemble-preparation and filter-burning steps is approximately 150 ms.

After opening the shutter in front of detector D_R, the RASE experiment is initiated by inverting the population. Rather than initially creating an absorption feature in either of the ground states addressed in the 4LE sequence, this RASE experiment begins with these two ground states empty, and a π-pulse at ω_{14} along path c to invert population. The process of population inversion also involves applying an electric field-gradient to broaden the gain feature, and additionally to reduce the gain per unit bandwidth along the entire c-d path so as to reduce superradiance effects.

Figure 7.3(c) depicts the inversion π-pulse at ω_{14} applied along path c. For the duration of this pulse the field gradient is switched on. Note that after many iterations of the entire pulse sequence, the ω_{14} π-pulse, in combination with the sweeping fields on during the first step, will select the specific frequency-class of ions, since any other frequency subgroups will eventually be pumped into one of their non-resonant ground states.[2]

Following the inversion π-pulse, the spontaneous emission is detected at D_A, before the rephasing π-pulses at ω_{24} and ω_{35} along paths c and d occur (fig 7.3(d)), and the rephased version of ASE that was detected at D_A will be emitted in the mode of D_R as given by the phase-matching condition, equation 7.1.

Finally, figure 7.3(e) highlights the important regions within the sample; the interaction region where the ASE will be rephased by the 4LE sequence, and the region where paths b and e overlap which constitutes the RASE frequency filter. An important point to note is that the population is inverted along the entire c-d path. Short of adding a 6$^{\text{th}}$ beam path, this is the only feasible option for inverting population in the interaction region, however there are problems associated with having a gain feature all the way along path c-d, as will be discussed in §7.2.5.

[2]This frequency-class selection is usually performed in a single step similar to the methods described in §6.4.

The APDs are both gated on for 20 μs following the inversion pulse for the ASE and, a longer period of 220 μs following the rephasing π-pulses, as shown in figure 7.4. Although there will only be RASE for the first $\sim$20 μs, the longer detection period is useful to quantify the level of spontaneous-emission background noise.

The maximum power available in beam-paths c and d was 7.6 mW and 19.8 mW respectively, such that the optimal pulse lengths (fwhm in amplitude) were 2.7 μs for π_{invert}, 1.2 μs for π_1 and 1.7 μs for π_2.

7.2.5 Creating the initial population inversion

The primary reason for beginning the RASE process by exciting from ground state $|1\rangle$ is to minimize the number of atoms that are addressed by the rephasing π-pulses which are *not* associated with any ASE or RASE.

For instance, one alternative would be to prepare an initial absorption feature in ground state $|2\rangle$, then apply a π-pulse at ω_{24} to generate the inversion. In the absence of any Rabi-frequency selection, only a fraction of these ions will be efficiently transferred up to state $|4\rangle$, and the majority remain the in the ground state (see §3.4.4). These ions which remain in the ground state will not contribute to any ASE signal, however following the rephasing π-pulses, they will be either absorbing or amplifying the propagating RASE signal, or worse still, spontaneously radiating into the RASE spatial mode.

Another inversion method that was investigated was to use a frequency-chirped pulse rather than a π-pulse. This is a very efficient way of inverting all ions indiscriminately of the Rabi-frequency that they see, however the rephasing π-pulses will still only efficiently rephase the fraction of these ions which experience pulse areas close to π. As such, there is a large ASE signal, but only a very small fraction of this is efficiently rephased.

Therefore, the optimal starting point for the RASE experiment, where the maximal number of ions initially inverted can contribute to ASE *and* RASE signals, is to use a π-pulse from the remaining ground state, $|1\rangle$.

The decision to invert the population from ground state $|1\rangle$ introduces a complication of its own due to the resulting spatial distribution of the gain feature. Given the beam geometry, as shown in figure 7.3(c), the π_{invert} pulse will invert population along the entire 20 mm length of the sample, while only 2.4 mm of this length overlaps with the detector spatial modes. Without any field gradient applied, the gain along the c-d spatial mode will be much larger than along path a-b as shown in figure 7.5(a). To quantify this, consider the transmission through an amplifying medium;

$$T = e^{\alpha L}$$

Where L is the path length, and α is dependent on the ion spectral density and other material properties. If $T = 2$ along the 2.4 mm path length of the overlap region, meaning any photon spontaneously emitted into this mode will stimulate on average one other photon event, then along the 20 mm length of the c-d path, $T = 320$. This will trigger significant stimulated emission into the c-d spatial mode. The inverted population will be rapidly depleted, effectively shortening the lifetime of the inversion feature[3] (see inset of figure 7.5). This means that the period of recording ASE must be shortened correspondingly to ensure that the ensemble is mostly in the ground state for the rephasing (or else adding noise to the RASE signal), and thus losing temporal modes for the entanglement generation.

A secondary problem that arises with this method of inversion, is that the π-pulse propagates through a significant optical depth before reaching the overlap region. As the pulse deposits energy, it becomes temporally stretched, and therefore spectrally narrowed (one interpretation of the 'area theorem' [94]). It will then only be an effective π-pulse for a much smaller spectral region than the input pulse length would suggest. This is a technical issue which could be solved with unlimited laser power. With the finite 20 mW of peak power available in the path c spatial mode, the length of π_{invert} was 2.7 μs, and the maximum achievable width of the gain feature was only $\sim$100 kHz. Since this lengthens the time-bins in which correlations can be expected, this narrow gain feature means that again the number of temporal-modes for generating ASE/RASE pairs per shot is reduced.

Reducing the number of temporal modes in which entangled photons can be generated is detrimental in this experiment for two reasons. The first is a practical issue: The lengthy preparation of the various spectral and spatial properties of the ensemble ($\sim$150 ms per shot) translates to long data acquisition periods ($\sim$hours) in order to draw statistically meaningful conclusions. The entire data-acquisition duration can be reduced in proportion with the number of temporal modes for RASE generation each shot.

The second issue with a limited number of temporal modes is conceptual in nature: One of the ways in which RASE differs from the DLCZ protocol is that it is multi-mode in time. This characteristic is inherited from the photon-echo upon which it is based. Therefore it would be desirable to achieve a regime where this temporal multi-mode nature could be explored.

These difficulties associated with a large gain, namely the lifetime-shortening and the narrow spectral width of the inversion, may be mitigated by reducing the gain along path c-d resulting from the inversion π-pulse.

One possible means of achieving this would be use an alternate sample geometry, such that only ions in the interaction region are inverted. This could be achieved for instance using two separate samples; one for generating RASE, and a second located along path a-b to implement the frequency filtering. Another method would be to introduce an additional independant beam path for applying the π_{invert} pulse. However the method chosen to effectively reduce the gain along path c-d in this experiment is to introduce a position-dependant Stark-shift. This represents the simplest solution to implement, requiring only

[3]This reduced lifetime was seen in modelling spontaneous emission from an ensemble of atoms in Appendix B. It is one characteristic of an effect known as superradiance - the borderline between incoherent and coherent emission. Superradiance was first examined theoretically by Dicke in 1954 [110], and is also discussed in texts [28, 94, 125].

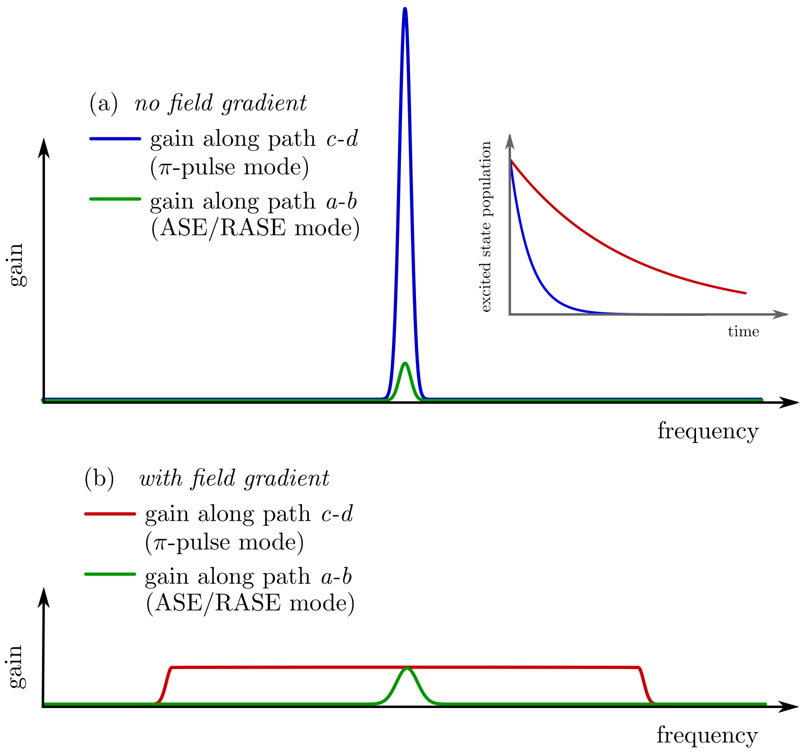

Figure 7.5: Illustrating how the initial population inversion in the RASE protocol can be altered by applying a field gradient during the inversion π-pulse. (a) Without the field gradient, probing the gain along the ASE/RASE spatial mode (path *a*-*b* in fig 7.2) is simply sampling a small fraction of a spectrally-identical, but much larger, gain along the entire length of the sample in the π-pulse mode (path *c*-*d*). (b) With the field gradient on, ions from a range of frequencies are Stark-shifted into the π_{invert} frequency window, such that when the gradient is switched off the gain feature is broader in frequency. In this case the gain along path *a*-*b* is similar in magnitude to the gain in the same frequency bandwidth along path *c*-*d*. Inset shows the lifetime-shortening associated with a large gain.

the addition of electrodes to the sample holder.

A quadrupole electrode arrangement (discussed in more detail in §7.2.6) results in a changing electric field along the length of the sample. This causes a position-dependent Stark-shift of the ions. For example, if ions at the front end of the sample experience a detuning of Δ, then at the back end the ions will shift by $-\Delta$. In the middle of the sample the field is zero, so there is no net shift in transition frequency. If the field gradient is switched on for the duration of π_{invert}, then there are ions throughout the sample that are Stark-shifted *into* the frequency window addressed by π_{invert}. When the field gradient is then switched off, the inverted ions are spread over a $\pm\Delta$ frequency window as shown in figure 7.5(b).

With the broadened gain profile, the gain along paths *c*-*d* and *a*-*b* is similar in the frequency window of interest. It was seen that the lifetime of the gain feature was not significantly shortened. Additionally, the spectral profile of the gain along the detection path *a*-*b* was seen to be broadened by a factor of $\sim$2 due to the field variation across the overlap region, which helps in narrowing the temporal width of the correlation between the photon pairs. These measurements will be discussed further in §7.3.

7.2.6 Sample holder and orientation

To apply the electric field gradient, the sample is surrounded by electrodes in a quadrupole arrangement, as shown in figure 7.6. There are also Delrin plates at the front and back with two apertures cut in each to assist with getting the two beam paths overlapping in a consistent manner, despite daily variations in beam pointing of the laser output. The front and back electrodes sit below these apertures, and the copper blocks along the side of the sample have slits cut through them to allow optical access through the sides.

The model in figure 7.6(a) was constructed using finite element analysis software, COMSOL.[4] The electric field was then numerically calculated, with the front and back electrodes charged to 15 V, and both side electrodes grounded. The results for the electric field along the $\hat{z}$-axis through the centre of the sample are shown in figure 7.6(b).

The sample itself is oriented such that the crystal *b*-axis is along the $\hat{z}$ direction, and the transition dipole axis is predominantly along the $\hat{y}$ direction.

7.3 Spectrum and lifetime of the inversion feature

The spectral properties of the entangled ASE-RASE photon pair are fundamentally defined by the spectrum of the gain on the ASE transition following the inversion π-pulse. The magnitude of the gain, in addition to affecting the emission rates via superradiance, also has an important but less obvious role in defining the entanglement dynamics[5]. For example, in the ratio of cross- to auto-correlation functions (R) the smaller the mean

[4] However the figure itself was generated using Autodesk Inventor.

[5] Although this depends on the specific metric being used to quantify the non-classical nature of the fields, as discussed in §5.

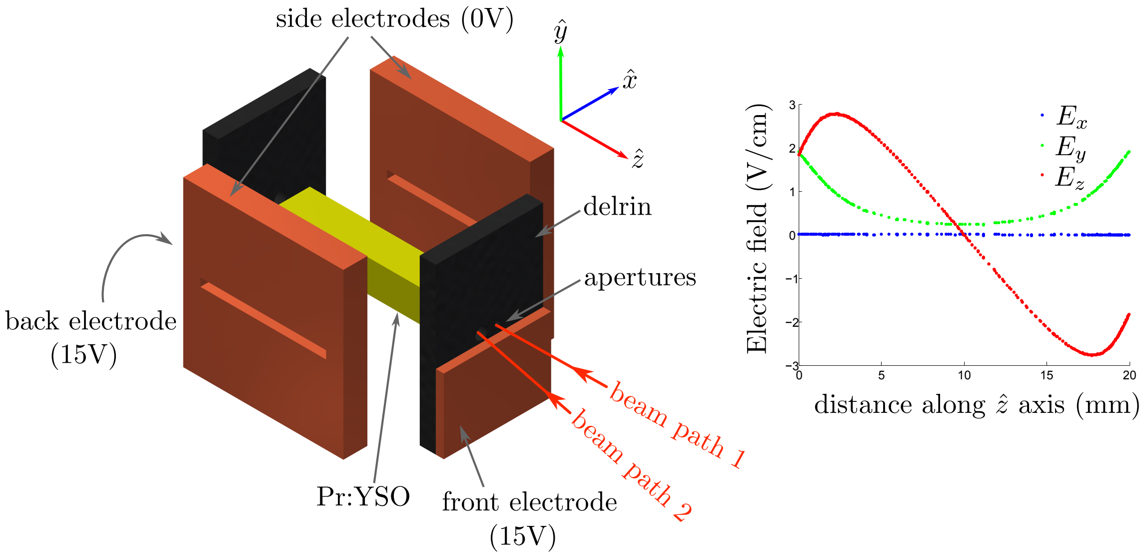

Figure 7.6: Model of the sample and copper-electrode arrangement. Not shown are the Teflon top and bottom plates which hold all the pieces in place. Also not visible is the back electrode, which is identical to the front, and is also charged to 15V. The side electrodes are both grounded. The graph shows the components of the electric field along the length of the sample (x=0, y=0, $-10 < z < 10$mm) calculated using finite-element analysis (using COMSOL).

number of excitations $\langle n \rangle$, the larger the violations of the Cauchy-Schwarz inequality (see equation 5.16).

The magnitude of the gain, specifically the value of $\langle n \rangle$, will be determined directly from the count histogram of the RASE data in a later section (§7.6). In the current section, measurements of the gain spectrum made using heterodyne detection will be presented. Also, results for the emission lifetime, measured (using single-photon counting) as a function of the applied electric field gradient will be discussed.

In determining the excited state lifetime, the spontaneous emission was measured by the APD D_A for the 150 μs following the inversion π-pulse. The results are shown in figure 7.7, for varying electrode voltages applied during π_{invert}. This confirms the effectiveness of the field gradient in reducing the superradiance along the high-gain path of c-d.

The lifetime of the optical transition for isolated ions is approximately 160 μs. However when there is no field gradient applied to the inverted ensemble, the decay time of the spontaneous emission into the path a-b mode is 14 μs (found by fitting an exponential to the first 15 μs). With the field gradient on, the decay time for this period is 98 μs. In this case, the superradiance effect is diminished, but still evident.

To characterize the spectral profile of the gain, the transmission of a short probe pulse along path a following the inversion pulse was recorded using heterodyne detection (this method was described in §3.4.2). The transmitted spectral amplitude, both with and without the field-gradient applied during the inversion, is shown in figure 7.8. The field gradient is seen to increase the fwhm of the gain feature from 83 kHz to 163 kHz, while reducing the intensity gain by a factor of 4.5. This broadening is consistent with

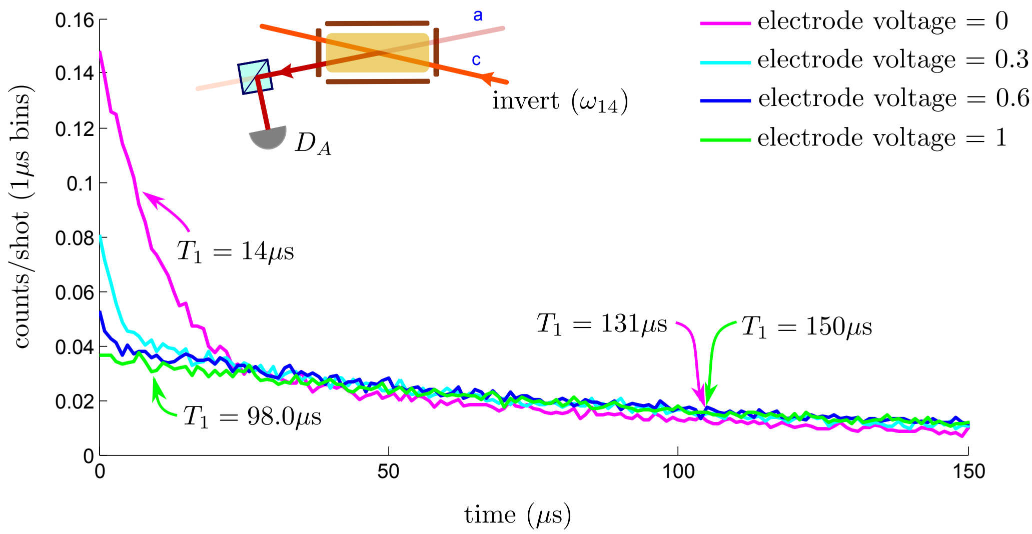

Figure 7.7: Quantifying the effect of the electric field distribution. This is a histogram of counts recorded in detector D_A following an inversion π-pulse along path c as shown in the inset. Each trace has a different electrode voltage applied during the π_{invert} pulse, and this voltage is given as a fraction of the maximum, which was 15 V. During the first 15 μs, the lifetime (T_1) of the emission with zero field-gradient is 14 μs; significantly faster than the transition lifetime of 160 μs. The same period with the maximum field gradient applied has a decay time of 98 μs. In the period 100-115 μs after the π_{invert} pulse, these decay rates slow to 131 μs and 150 μs for no field, and maximum field respectively. The total number of shots per trace was 15000.

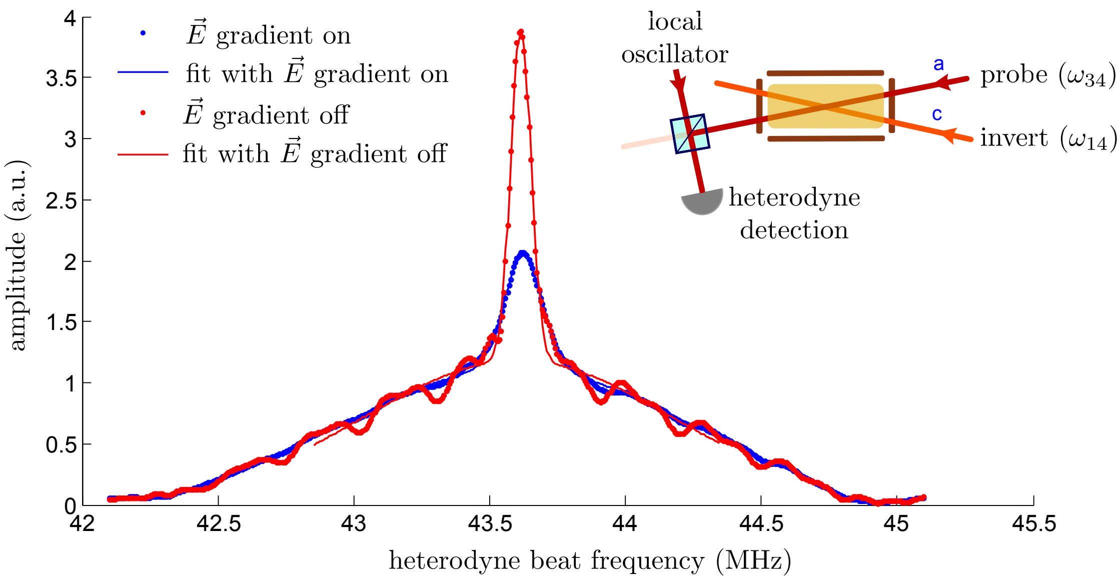

Figure 7.8: Gain with and without electric field gradient applied. This plot shows the measured transmission of a weak probe pulse along path a, occurring 5 μs after π_{invert} and detected using heterodyne for two cases; where the field gradient during π_{invert} was on or not. Without the field-gradient on, the gain along the interaction region in terms of intensity was 15.2 (or an inverted optical-depth parameter αL of 2.72), and the feature width was 83 kHz (gaussian fwhm). With the field-gradient on, the gain feature width is broadened to 163 kHz, and the amplification factor reduced to 3.4 ($\alpha L = 1.22$). The beat frequency between local oscillator and the centre of the gain feature was 43.6 MHz, and the probe pulse fwhm was 2.3 MHz, and each trace is averaged over 50 shots.

field-gradient	maximum delay time	photon width	temporal modes per shot
ON	22 μs	6.1 μs	3.6
OFF	3.1 μs	12 μs	0.26

Table 7.1: Temporal-modes available for photon pair generation per shot

an interaction length of 2.5 mm using the reported value of the Stark shift in Pr:YSO of 112 kHz/(V.cm^{-1}), predominantly along the $\hat{z}$-axis, from reference [91].

To summarize the effect of the field gradient, consider as a figure-of-merit the number of temporal-modes in a single shot in which ASE and RASE photon pairs can be expected. This is determined as the detection time window divided by the length of a single temporal-mode. Firstly, the delay time between the inversion and the rephasing pulses (the period for detecting ASE) is chosen as the time when the emission is reduced by 20% from its maximum value. This is a reasonable threshold to ensure the rephasing process is not noisy.[6] The temporal width of an ASE/RASE wave-packet (i.e. the width of a single temporal-mode) is approximated as the inverse spectral-width of the gain feature. The resulting number of temporal modes available per shot for generating photon pairs is 3.6 with the field-gradient applied, or 0.26 without (see table 7.1).

Although this figure-of-merit neglects the four-level echo coherence time (this will be reported in §7.5), it is still apparent that utilizing the Stark-shift will significantly increase the number of temporal modes per shot. In turn, this increases the number of modes in which high-fidelity ASE and RASE photon pairs can be generated, improving the overall efficiency of the process.

7.4 Background noise and frequency filtering

As shown in Section 5.6, a relatively small amount of *incoherent* noise (which has thermal statistics - see Appendix D) can completely mask any non-classical correlation between ASE and RASE photons. Also, the RASE signal itself will be reduced due to the low efficiency of the rephasing process. Therefore maintaining a reasonable signal to noise radio in the RASE mode will be a priority for achieving any measurable degree of quantum correlation between the photon pairs.

The coherent FID noise following the rephasing π-pulses is already reduced to a negligible level by the counter-propagating beam geometry and the 10° angle between detector and π-pulse modes. And the bandpass interference-filters in front of each detector blocks most of the emission to intermediate crystal-field levels - the 'red' emission.[7]

What remains as a potential noise source is spontaneous emission from either excited

[6] Recall that the 'quietness' of the RASE process relies on having most ions in the ground state when the rephasing is happening.

[7] See figure F.3 for the measured emission spectra of Pr^{3+} : Y_2SiO_5, and figure 2.1 for the energy level diagram.

state ($|4\rangle$ or $|5\rangle$) to either ground state ($|2\rangle$ or $|3\rangle$).[8] Three of these transitions are distinguishable from RASE in frequency and can be filtered out. Spontaneous emission noise at the RASE frequency however poses a more fundamental problem, which will be discussed in further detail in §7.9.1.

The best option for frequency filtering with the required MHz discrimination is to use hole-burning properties of the sample itself, particularly for filtering weak fields which are unlikely to cause hole-burning themselves. This spectral-tailoring is the purpose of beam path e, and where this beam overlaps with the detector mode (path a-b) will form the filter region. This is marked in figure 7.3(e).

The atomic population is initially transferred to the excited state $|4\rangle$, and the ions which radiatively decay down to ground state $|3\rangle$ will be recorded as ASE. The quantum efficiency of these *direct* transitions however is quite low, with branching ratios of $\sim$2%. Thus, there will be many ions that end up in ground state $|3\rangle$ which decay via intermediate crystal-field levels. Also, since the transition strengths are comparable, the ground state $|2\rangle$ is similarly populated. Therefore, following the rephasing π-pulses, all of the ions which are excited from $|2\rangle \rightarrow |4\rangle$, and the 'incoherent' ions (which weren't associated with any ASE events) which are excited from $|3\rangle \rightarrow |5\rangle$, can potentially decay back to any ground state and spontaneously radiate into the RASE detector mode. Also, there may be ions which remain in the excited state $|4\rangle$ throughout the ASE-detecting period, and aren't efficiently transferred to the ground state by the π-pulses. Again, these ions spontaneously radiating constitutes another noise source.

Of all of these potential noise transition-pathways which can be resolved in frequency from the RASE, the dominant component will be at the ASE frequency, since there is a transmission window explicitly hole-burnt before every shot along the detector path for the purposes of detecting ASE at D_A. It was seen however, that scatter from the fields at the π-pulse frequencies was also sufficient to eventually (after a few minutes) hole-burn through the sample.[9]

Figure 7.9(b) shows all the frequencies at which noise is expected, both in the ASE and RASE temporal periods. The aim for the RASE filter region is to burn-back some ion population which will absorb the three frequencies shown, namely ω_{35}, ω_{24} and the ASE frequency ω_{34}. Since the experiment only utilizes one single frequency-subgroup of ions,[10] the other subgroups can be used to provide this absorption. Figure 7.9(a) shows the frequencies that were applied to burnback population that would absorb emission to ground state $|3\rangle$ in the experiment subgroup, and similarly (b) shows the subgroups that were used to provide absorption of emission to ground state $|2\rangle$.

The required frequencies for the filter-burning process were generated using a double-pass AOM configuration, driven sequentially by four fixed-frequency RF channels. This

[8] There is also spontaneous emission from excited states $|4\rangle$ and $|5\rangle$ to ground state $|1\rangle$, however these transitions occur at the $\sim$5% level, while 95% of the population ends up in states $|2\rangle$ or $|3\rangle$.

[9] One downside of having a laser stabilized to <1 kHz linewidth.

[10] Recall that for any single applied optical frequency there will be 9 resonant frequency-subgroups. This is because the transition is inhomogeneously broadened.

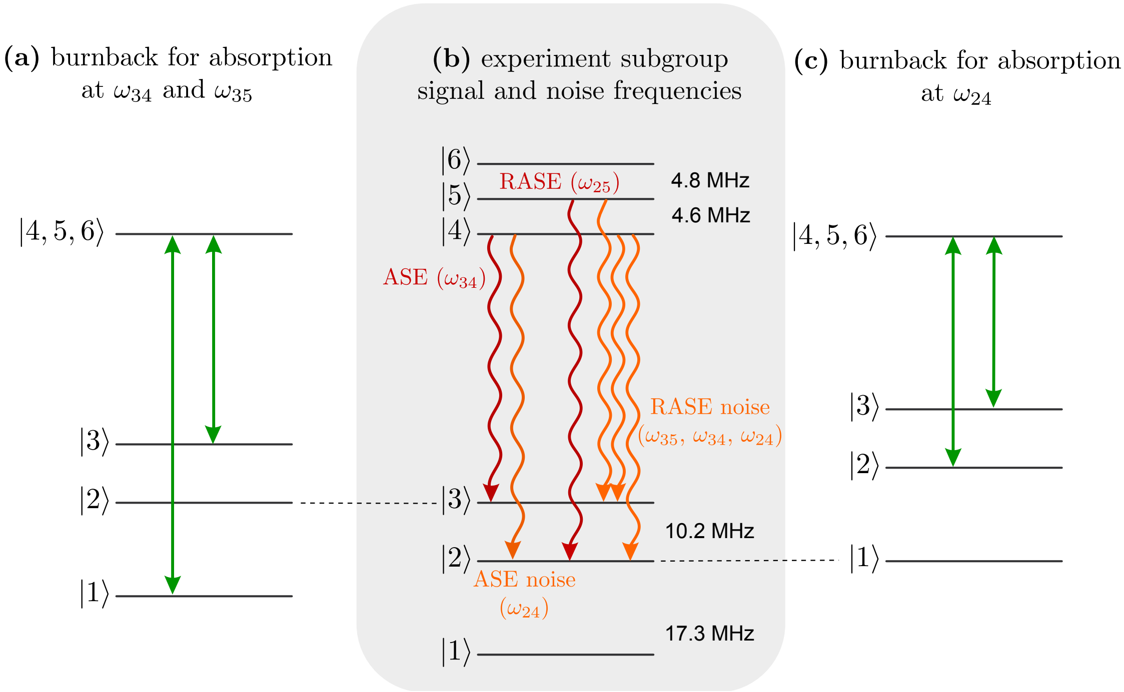

Figure 7.9: Signal and noise emission frequencies for RASE experiment in Pr:YSO. The centre figure represents the single subgroup of ions in which RASE is generated. The noise frequencies are marked in orange, and the ASE/RASE signal frequencies are marked in red. Subfigures (a) and (c) show the *different* subgroups of ions used for burning-back population to absorptively filter out the spontaneous emission noise in the RASE mode.

beam was then passed through an EOM, which was driven by a signal sweeping in frequency from 20 kHz to 200 kHz (at a rate of 1 kHz). This chirps the optical frequency from its centre value through to $\pm$200 kHz, and broadens the otherwise narrow absorption features that are burnt back. While the four burn-back fields along path e were applied iteratively, the field burning the transmission window for the RASE along path a-b remained on to ensure that the RASE itself will not be absorbed (as depicted in figure 7.3(b)).

To quantify the absorption of this filter at the frequencies of interest, the transmission of four probe pulses (each 1.7 MHz fwhm, and centred at ω_{25}, ω_{24}, ω_{35} and ω_{34}) was measured using heterodyne detection with and without the filter-burning field switched on. The results are shown in figure 7.10. The most important features to note are that the RASE transmission is unaffected when the filter-burnback is on, and that the ASE frequency is sufficiently absorbed.

Although the absorption at the π-pulse frequencies ω_{24} and ω_{35} doesn't appear to be very high in figure 7.10, this is because the probe pulses themselves are hole-burning. During the actual experiment, the only light at this frequency in this spatial region is due to scatter from path c-d. Thus some slight burning back at these frequencies should be sufficient to maintain a strong absorption. Typically the attenuation averaged over the 160 kHz wide region centred at the ASE frequency was found to be 26 dB.

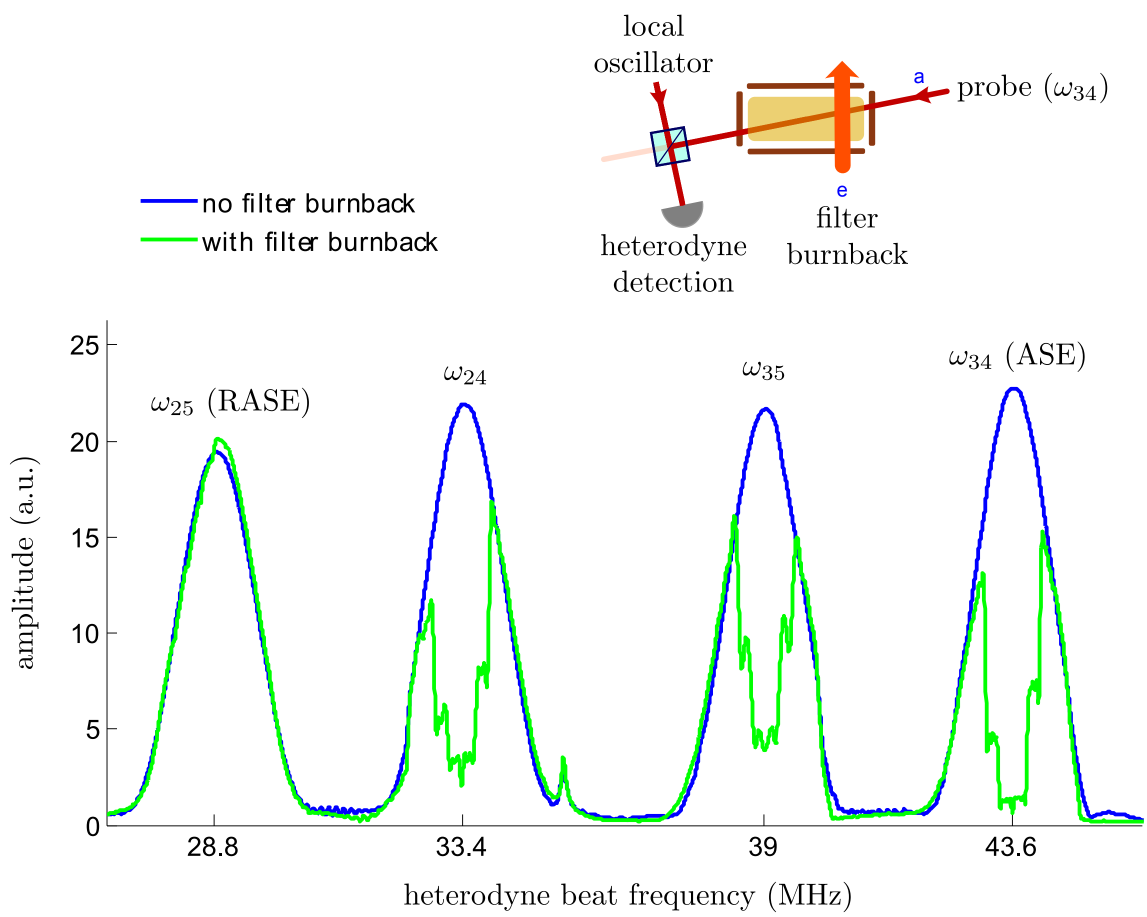

Figure 7.10: Quantifying the transmission of the hole-burnt frequency filter. This figure shows the amplitude spectrum of transmitted probe pulses with and without the filter-burning field along path e switched on, averaged over 100 shots. The structure of the burnt-back absorption profiles reflects the properties of the beam after transmission through the EOM; the light is mostly shifted to the first order region, ±200 kHz from the centre frequency, with some light also shifted to the second order. The key transmission features of the filter are that the RASE frequency is unaffected, and that the ASE frequency is heavily absorbed (since there is a spectral hole burnt every shot in this spatial mode). Typically there was 26 dB of absorption in the ASE frequency window.

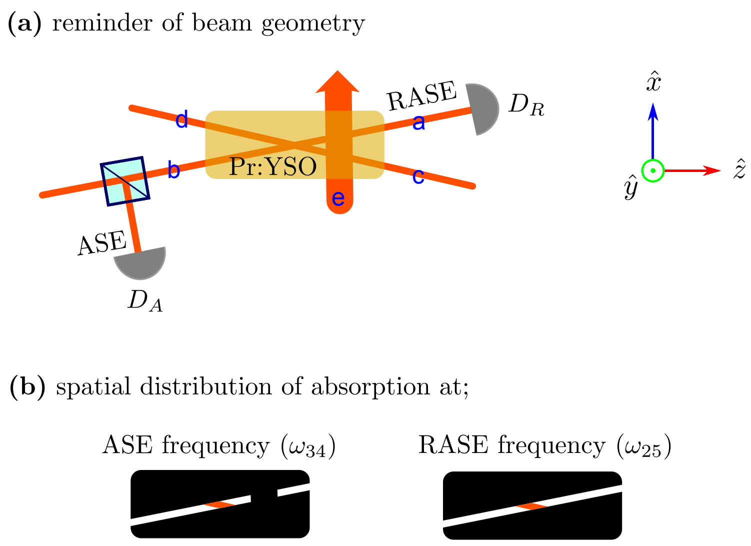

Figure 7.11: Summary of the sample absorption properties, showing the *spatial* distribution for the two main *spectral* regions of interest; the ASE and RASE frequencies. In subfigure (b), the regions of the sample coloured black represents absorption at the given frequency, while white indications transparency. The interaction region is marked in orange.

In summary, frequency filtering for the detection of RASE was achieved by tailoring the absorption spectrum of the sample in the region where beam path e intersects with the detection modes of path a-b. Figure 7.11 depicts the spatially varying spectral properties of the sample for the ASE and RASE frequencies; in the regions coloured black the sample is opaque, and in the white regions the sample is transparent.

7.5 Echo efficiency and coherence time

This section describes measurements of the four-level echo rephasing efficiency, and the coherence time. Similar measurements were made in the initial characterisation of the 4LE in the co-propagating geometry, described in Chapter 6. In the counter-propagating geometry however, the rephasing efficiency is reduced due to difficulties in spatial mode matching. A further difference in this geometry is that the region of the sample generating the echo (2.4 mm long) has been much reduced compared to prior measurements which used the entire 20 mm length. The cancellation of the residual magnetic-field (described in §6.4) is re-calibrated accordingly, and this extends the coherence time.

To characterize the echo, an absorption feature was created in ground state $|2\rangle$. An input pulse along path b at ω_{25} can then be rephased at ω_{34} and detected at D_A. Figure 7.12 is a plot of the echo intensity as a function of the delay time, measured using heterodyne detection.

As described in §6.4, the residual magnetic field at the sample was nulled by current-carrying coils placed outside the cryostat. Although the echo amplitude decay is not purely exponential, an approximate coherence time for the relevant time window for the

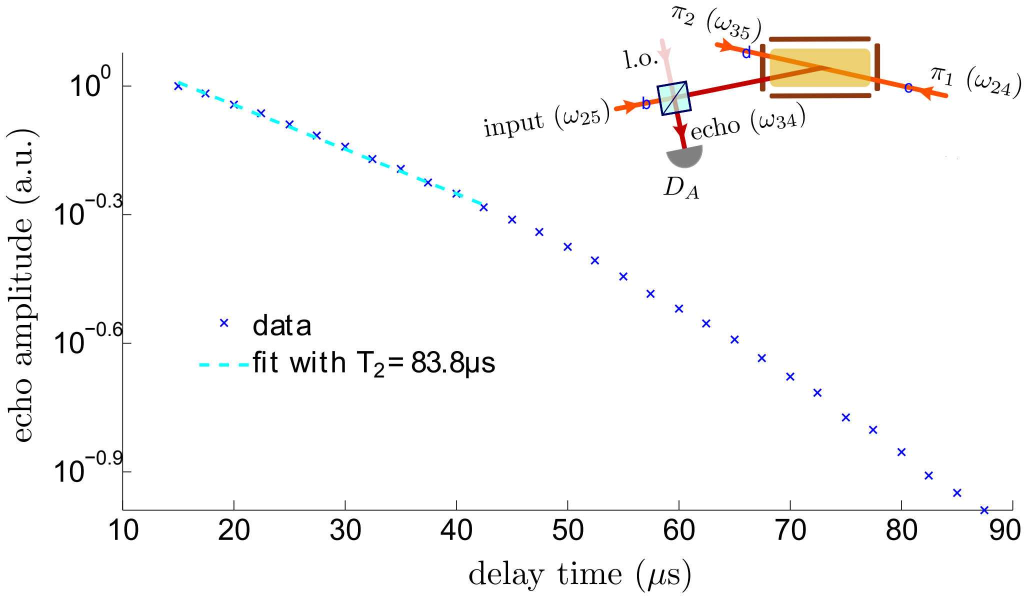

Figure 7.12: Measured four-level echo amplitude, in the counter-propagating geometry, determined with heterodyne detection, as a function of delay time between π-pulses and the echo. Although the decay is not an exponential, the early stage can be approximated with an exponential fit from which a coherence time T_2 of 83.8 μs is obtained. Each point represents the average of ≈35 shots.

RASE experiment was obtained by fitting[11] the decay in the initial 40 μs period after the π-pulse, giving the value of T_2=84 μs.

The reason for this improvement from previous measurements (from figure 6.4, T_2=34 μs) is twofold. Firstly, only the initial decay region is being considered here, and the decay rate is increasing at longer times. The average decay over the full 100 μs domain is 63 μs. The second cause for an improved T_2 is due to the smaller interaction region. The cancellation of the residual magnetic field is achieved to a higher degree of accuracy across the 2.4 mm length, compared to the full 20 mm of the sample length used in chapter 6.

Heterodyne measurements of the echo were also used to confirm that the full width of the gain feature was being rephased. In this case, the ensemble began in the excited state $|4\rangle$. A weak input pulse at ω_{24} was used to generate some coherence to be rephased by the π-pulses. The spectral fwhm of the echo was found to be 164 kHz, matching the width of the initial gain feature.

The counter-propagating geometry used here will be less efficient than the co-propagating version due to imperfect overlap of the different beam paths. This efficiency was quantified in two different settings; either starting from an absorption feature (population in state $|2\rangle$) or a gain feature (population in state $|4\rangle$). In both cases, a weak (mean number of photons ≈500) input pulse along path b at frequency ω_{24} was rephased as an echo in the opposite direction. For these measurements the APD detectors were used. The signal histograms are plotted in figure 7.13 for the case where there was an initial absorption

[11]The fit is of the form: echo amplitude = $\exp(-t/T_2)$, where t is the time from input pulse to echo.

feature. The transmitted input pulse when there was no absorption/gain feature was also measured to give the magnitude of the input pulse.

Integration over the peaks in figure 7.13 gives the average detected count rates as

$$\begin{aligned}\text{Input pulse magnitude} &= 53.9 \quad \text{photons/shot}\\ \text{Input pulse transmitted} &= 21.5 \quad \text{photons/shot}\\ \text{Echo} &= 1.41 \quad \text{photons/shot}\end{aligned}$$

The optical depth of the initial absorption feature was determined to be $\alpha L \approx 0.92$. In the context of the RASE experiment where there is no input-pulse, the relevant efficiency measurement here is the fraction of the photons that are actually *absorbed* by the ensemble (a measure of the degree of ensemble excitation) compared to the number of photons in the echo. For the above count rates, this is 4.3%. Scaling the counts according to the different transmission efficiencies to the two detectors, and also extrapolating back to zero delay time using the previously measured T_2, gives an estimate of the echo efficiency as 25.9%±0.7%.

For the similar experiment with initial population in the excited state, the efficiency is determined as the echo intensity relative to the number of photons that are *stimulated* by the input pulse (this represents the degree of de-excitation of the ensemble). The experiment was repeated for three different magnitudes of the gain. The efficiency (assuming zero delay time, and correcting for transmission losses) was found to be 50.1%±0.3%, 50.3%±0.7% and 47.5%±1% for $\alpha L \approx -0.62$, -0.30 and -0.15 respectively. Somewhat surprisingly, there is very little difference in efficiency across the range of gain settings measured. However since the echo was only generated for a single delay time in each gain regime, a variation in the lifetime (and therefore coherence time), cannot be ruled out.

In previous measurements characterizing the four-level echo with a co-linear geometry the rephasing efficiency (of only the absorbed part of the input pulse) would be close to 100%. However a lower efficiency is to be expected in the counter-propagating geometry given the difficulty in perfectly mode-matching all beams involved. Since the π-pulse beams were entirely free-space modes, it was not possible to directly quantify the mode overlap. Also, the two sets of efficiency measurements, starting from either an absorbing or amplifying medium, were performed on different days and thus beam overlaps could have been quite different. The cause for the disparity in efficiencies can be attributed therefore to either of two effects; differing propagation properties between the rephasing process starting from a gain or absorption feature, or simply a difference in the quality of beam mode-matching.

In summary, the measurements in this section confirm that the four-level echo can be generated with counter-propagating π-pulses, and starting from either an absorbing or an amplifying medium. Echo efficiencies in the range of 25-50% were achieved.

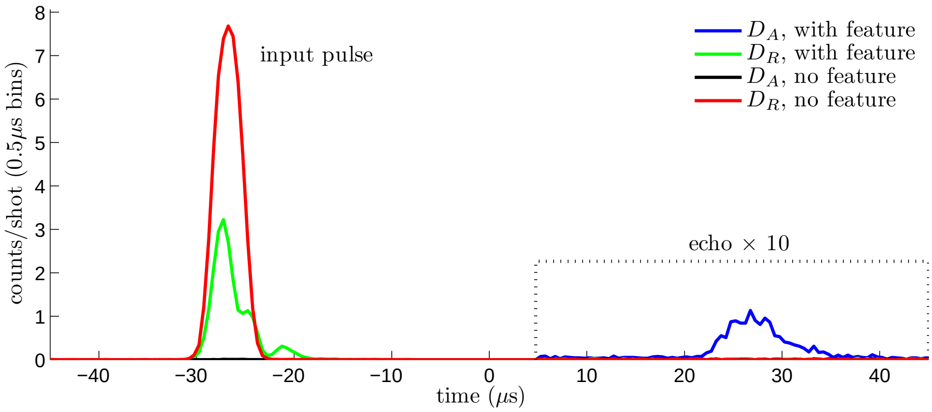

Figure 7.13: Measurement of echo efficiency, using single photon counting detection. The plot shows the average counts recorded across 1000 shots in 0.5 μs time-bins by both APDs D_A and D_R during a four-level echo sequence starting from an absorbing feature in ground state $|2\rangle$ of optical depth 0.92.

7.6 RASE experiment results: count histogram

Having quantified the properties of the inversion feature, the filter transmission, and the echo efficiency and coherence time, the data acquisition for the RASE experiment itself was performed. In this section, we examine the histogram of counts averaged over the entire data acquisition, and estimate the ASE/RASE signal and background noise levels in each detection channel.

As shown schematically in figure 7.4, during the RASE experiment the APDs are gated on for 20 μs following the inversion π-pulse, and again following the rephasing π-pulses for 220 μs. The output of the APD detectors are time-stamped relative to a 'start' pulse using an FPGA (see §F.4) with ∼20 ns resolution. A total of 10^5 shots were recorded over approximately 5 hours.

Figure 7.14 shows the count rates measured in 1 μs time-bins The evidence that there is indeed rephased spontaneous emission is in the blue trace of figure 7.14, where there is a peak following the rephasing pulse which is comparable in width to the ASE emission window.

To partially negate the 'afterpulsing' effect of the detectors, any counts that occur less than 80 ns after another count are neglected in the data-processing. Afterpulsing occurs when an electron is captured within impurities in the semiconductor during an avalanche, and then is released after the detector dead-time and triggers another pulse. Silicon APDs also have a tendency to emit light from the avalanche region as a photon is detected. Since the experimental setup here involves a pair of APDs facing each other, the data was processed to check for evidence of this. It was found that the probability of getting a count at the same time in both detectors was at the level expected assuming the counts were independent, and therefore this effect was not visible.

The background count level (due to detector dark counts and scattered light) was

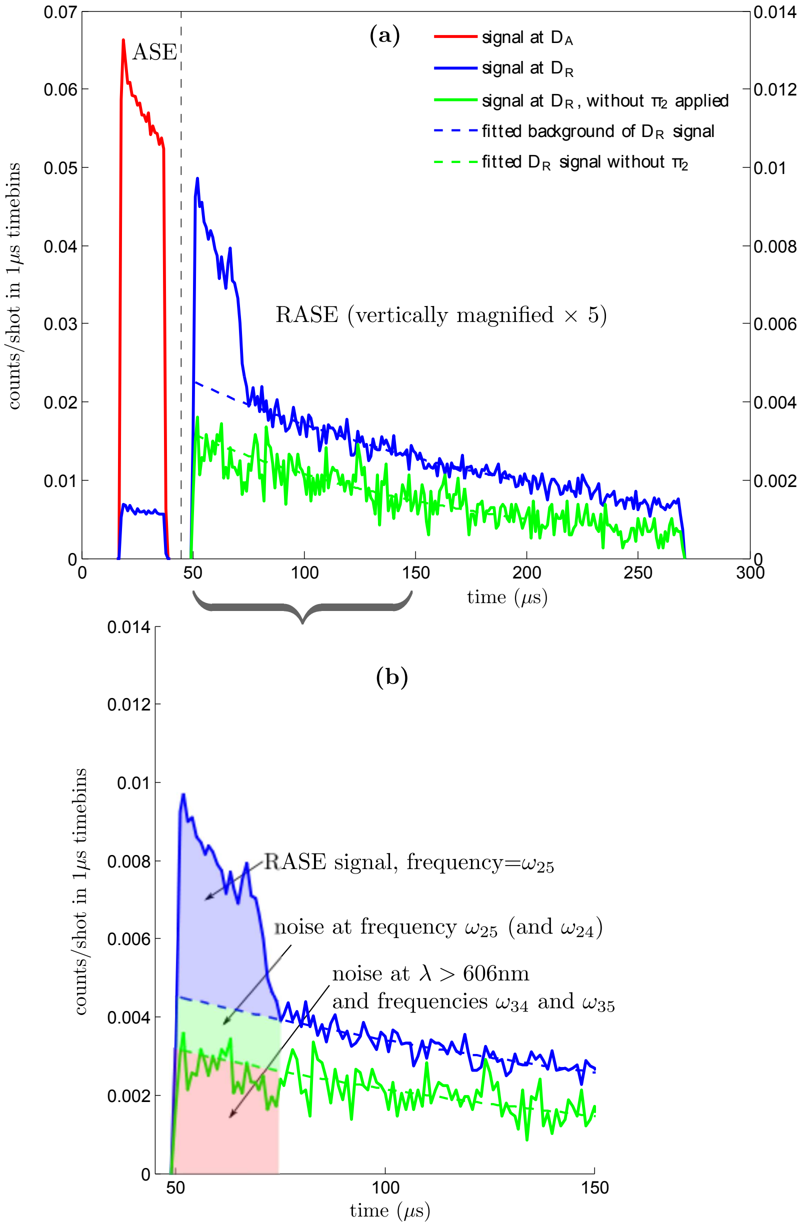

Figure 7.14: Histogram of photon counts in RASE experiment. The count histograms for the full RASE sequence is compared to the count histogram without the second rephasing π pulse applied, such that the signal and noise levels can be determined. (a) Evidence of RASE is visible in the D_R signal (blue trace) in the 50-75 μs period. The width of the RASE signal matches the ASE temporal region, and also this signal disappears if the second π-pulse is omitted (green trace). Note that the histograms from 50μs onwards have been scaled by a factor of 5, and the scale is shown on the right hand y-axis. Both detectors were gated on during the ASE and RASE temporal periods, however the D_A histogram (red trace) is only plotted before the rephasing π-pulses for clarity. (b) Close up of the RASE temporal region, where the noise contributions of different spectral modes are shaded.

deemed negligible in contrast to the signal count rates. This was determined by recording 4×10^3 shots where all of the π-pulses were switched off, while the feature preparation remained unchanged. The average background count rate across the entire detection period for detector D_A was $2.0\times10^{-4}\mu s^{-1}$ and was $4.3\times10^{-4}\mu s^{-1}$ for D_R. The detector dark count rates of $\sim 10^{-4}\mu s^{-1}$ accounts for a significant fraction of the background level. The higher count rate in D_R was attributed to laser leakage through the AOM and fibre system of path b.

The quality of the filtering was determined by comparing the intensities measured at the two detectors during the ASE temporal region.[12] As measured using heterodyne detection (see §7.4), there is 26 dB of absorption at the ASE frequency in the mode detected by D_R. However, the photon-counting measurements indicate an intensity difference of just 12 dB between the D_A and D_R counts (after accounting for the difference in transmission efficiencies along paths a and b). This is thought to be due to leakage through the bandpass filter of the red emission. Given the 3% branching ratio of the $^3H_4 \rightarrow {}^1D_2$ transition,[13] the counts detected in D_R indicates an overall attenuation for these longer wavelengths of 27 dB.

With 27 dB of attenuation, the emission at longer wavelengths accounts for less than 0.2% of the ASE signal. Therefore the ratio of the signal (at ω_{34}) to noise (all other frequencies) of the ASE detection is largely determined by the oscillator strength of the ASE transition. Thus the SNR is calculated as $\sqrt{0.55/0.45} = 1.1$

Using this SNR, and the transmission efficiency of the ASE mode, the actual signal level can be determined. The peak count rate occurs during the ASE period at approximately 0.06 μs^{-1}. Using the 5.8% probability of emission from the sample being detected in this mode, the approximate temporal width[14] of the photon wave-packets of 6 μs, and the SNR, then there is on average 3 indistinguishable ASE photons emitted in a single temporal mode each shot.

There is a large level of background emission from the sample itself when the full pulse sequence is applied. The signature of this emission, which is not associated with any ASE or RASE, is seen at long times ($> 80\mu$s) in figure 7.14. Fitting this temporal period with an exponential decay, and then extrapolating backwards (dashed blue trace), the signal to noise ratio of the RASE signal is estimated as 0.9.

As discussed in §7.4, the background noise consists of a number of different emission frequencies. To determine the frequency constitution of the noise, a series of shots was recorded where all pulses were on *except* for the second π-pulse at ω_{35}. This histogram is

[12]To clarify, the "ASE temporal region" refers to the period *before* the rephasing π-pulses (<45 μs in figure 7.14). Similarly, the terminology "RASE temporal region" will be used to imply the region *after* the rephasing pulses.

[13]The remaining 97% of emission is due to transitions to the ground state via intermediate crystal field levels (see energy level diagram figure 2.1). The largest emission peak is at 612 nm, and there is emission comparable to the 606 nm level out to 640 nm (see figure F.3 for a plot of the measured emission spectrum).

[14]The gain feature was measured as being 160 kHz (gaussian fwhm) in terms of field amplitude (see fig.7.8), so this corresponds to a width of roughly $\Delta\nu$=110 kHz in intensity. The Fourier-transformation of a gaussian (see for example [121]) gives a temporal fwhm $\approx 0.66/\Delta\nu = 6$ μs of the photons.

the green trace in figure 7.14. In this case, having only done the inversion at ω_{14}, and the first π-pulse at ω_{24}, the only emission detected during the RASE temporal period will be from level $|4\rangle$ to either of the ground states (which is attenuated by the sample-absorption filter) and any red emission (which is attenuated by the bandpass interference filter). This contribution to the noise, shaded red in figure 7.14(b), is in principle possible to filter out.

There is however a significant level of *additional* noise when the full rephasing pulse sequence is applied (the region shaded green in 7.14(b)). This component of the emission is problematic as this represents noise primarily at the RASE frequency itself, ω_{25}.

This background emission at the RASE frequency is due to atomic population that *indirectly* ended up in ground state $|2\rangle$, and once excited by the second π-pulse, spontaneously radiated into the RASE detection mode. In the current experimental configuration, this represents an unavoidable noise floor. With the measured count rates, assuming all of the emission at other frequencies is filtered out, a maximum signal to noise ratio of 2.8 may be expected. Methods to reduce this noise, which is indistinguishable from the RASE signal, will be discussed further in §7.9.1.

Finally, the rephasing efficiency can be estimated by comparing the ASE and RASE signal levels at time-symmetric points before and after the rephasing pulses. This indicates an echo efficiency of 10.2% at a point 8.9 μs from the focal point of the rephasing (which is determined using the cross-correlation in §7.7.1), and extrapolating back to zero delay time with the previously measured coherence time (see fig.7.12) the peak echo efficiency is 16%. This is lower than found in previous measurements using an input pulse, where the efficiency was in the range 25-50% (from §7.5). The reduced efficiency could be due to sub-optimal mode-matching between the two rephasing pulses caused by drifts in beam pointing during the 5 hour recording period. However it is also possible that there is a fundamental difference between an echo of an input pulse and an echo of spontaneous emission.

A summary of the important parameters determined from the histogram of counts shown in figure 7.14 is as follows;

$$
\begin{aligned}
\text{ASE photons emitted} &= 0.5 \quad \text{per } \mu\text{s per shot} \\
\text{ASE signal-to-noise} &= 1.1 \\
\text{RASE photons emitted} &= 0.04 \quad \text{per } \mu\text{s per shot} \\
\text{RASE signal-to-noise} &= 0.9 \\
\text{rephasing efficiency} &= 16\% \\
\textit{RASE signal-to-noise, with perfect frequency filtering} &= 2.8 \\
\textit{RASE signal-to-noise, without the hole-burnt filter} &= 0.3
\end{aligned}
$$

7.7 RASE experiment results: correlations between ASE and RASE

As discussed in Appendix D, the second-order normalized correlation function, $g^{(2)}$, is a useful way of quantifying correlations between two different modes. Evidence of a correlation between ASE and RASE modes provides a stronger confirmation (compared to the signature in the count histograms) that rephasing is actually happening. Also the correlation properties can be used to give some insight into the statistical nature of the process, including the affect of noise, losses, and mode-mismatch. Perhaps most importantly, the $g^{(2)}$ function can distinguish between 'classical' and 'quantum' correlations.[15]

7.7.1 Cross-correlations

If the ASE is being rephased, then the value of the (normalized) second-order cross-correlation function between the ASE and RASE modes, evaluated for points in time that are equidistant from the midpoint of the two π-pulses, will be greater than 1. If a slight time lag between the modes is introduced, the correlation will reduce. For completely independent signals, the cross-correlation is equal to 1. In this experiment, the width of the correlation peak yields information about the bandwidth of the RASE, while the magnitude indicates the degree of correlation.

The centre-point in time about which the rephasing occurs is labelled as t_0 (shown in figure 7.15). In evaluating the cross-correlation between the ASE and RASE modes, a 'coincidence' occurs when a photon in the ASE detector is recorded at time t **before** t_0, and then the RASE detector registers a photon at time t **after** t_0 in the same shot.

The time t_0, is not precisely known.[16] Therefore, there are two time parameters used to initially examine the cross-correlation; the time t_0, and the lag between 'coincidences' τ (defined in figure 7.15).

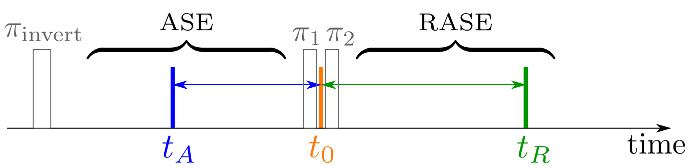

'lag time' $\tau = (t_R - t_0) - (t_0 - t_A)$

Figure 7.15: Definitions of the various times involved in determining the cross-correlation. The important parameters are t_0 (the time point about which cross-correlations are calculated) and τ (the lag time).

[15]The term 'non-classical' (or 'quantum') correlation implies a degree of correlation that can not exist according to a *semiclassical* model of the atom-field dynamics.

[16]This is due, for instance, to the delay between sending RF to the AOM, and the beam actually being diffracted. This varies on a day to day basis depending on where the laser beam passes through the AOM, and was not directly measured.

The function $g^{(2)}_{a,r}(\tau)$ (where subscripts a, r denotes the correlation is between ASE and RASE modes) was calculated for each value of t_0 using equation E.7. A total time window of duration $T = 14$ μs was chosen for the cross-correlation calculation, and the time lag variation is $-9\mu s \leq \tau \leq 9\mu s$. The resulting map of $g^{(2)}_{a,r}(t_0, \tau)$ is shown in figure 7.16(a). The vertical axis represents an effective shift of the time t_0. A negative offset defines a shift in the ASE time-window away from t_0, and a positive offset indicates a shift in the RASE window away from t_0. As expected for a photon-echo process, the correlation peak shifts in accordance with the offset. This is highlighted by the diagonal line along offset=τ.

Figure 7.16(b) shows the measured $g^{(2)}_{a,r}(\tau)$ for zero offset, which is a slice from the surface in (a) along the blue dashed line. A fit of the form given in equation D.17, with the amplitude of the gaussian component also allowed to vary, gives a coherence time $t_c = 4.4{\pm}0.8$ μs, and a peak amplitude of 1.21±0.03. Also shown is the average correlation between ASE and RASE modes across 10 adjacent shots (denominator of equation E.9). The mean cross-shot correlation value is 1.004, meaning that there is very little classical correlation from shot to shot. Thus each trial can be considered independent.

The expected peak value of $g^{(2)}_{a,r}$ in relation to the number of ASE/RASE photons can be determined using the model described in Chapter 4.

$$g^{(2)}_{a,r} = \frac{\langle N_a N_r \rangle}{\langle N_a \rangle \langle N_r \rangle} = 1 + \frac{T_a T_r \langle n \rangle (\langle n \rangle + 1)}{(T_a \langle n \rangle + \langle m_a \rangle)(T_r \langle n \rangle + \langle m_r \rangle)} \tag{7.2}$$

where subscripts a and r label ASE and RASE properties respectively. $N_{a,r}$ is the total number of photons detected, T_a is the transmission and detection efficiency along the ASE mode path, and similarly T_r is the net efficiency for the RASE mode, which includes the rephasing efficiency. $m_{a,r}$ are the background noise counts, and $\langle n \rangle$ is the mean number of ASE events emitted per shot.

As discussed in §7.6, the majority of the background noise is emission from the sample itself. Since this noise will scale similarly to the signal with respect to the transmission/detection efficiencies, it makes sense to re-write the noise parameters in terms of what gets emitted ($\langle M \rangle$) rather than what gets detected ($\langle m \rangle$), i.e. $\langle m_{a,r} \rangle = T_{a,r} \langle M_{a,r} \rangle$.

An additional experimental imperfection to consider is the mode-mismatch between ASE and RASE spatial modes. For a coupling efficiency of η_c, the number of correlated photon pairs generated reduces to $\eta_c \langle n \rangle$. Note that the remaining $(1 - \eta_c) \langle n \rangle$ counts constitute an additional noise term. Incorporating these parameters the expression for $g^{(2)}_{a,r}$ becomes;

$$g^{(2)}_{a,r} = 1 + \frac{\eta_c \langle n \rangle (\eta_c \langle n \rangle + 1)}{[\langle n \rangle + \langle M_a \rangle]\,[\langle n \rangle + \langle M_r \rangle]} \tag{7.3}$$

Note that the transmission loss terms $T_{a,r}$ have been cancelled out. The coupling efficiency prior to running the experiment was measured as $\eta_c = 50\%{\pm}5\%$. The peak value of $g^{(2)}_{a,r}$ according to equation 7.3 is 1.31, calculated using the average excitations per shot per μs

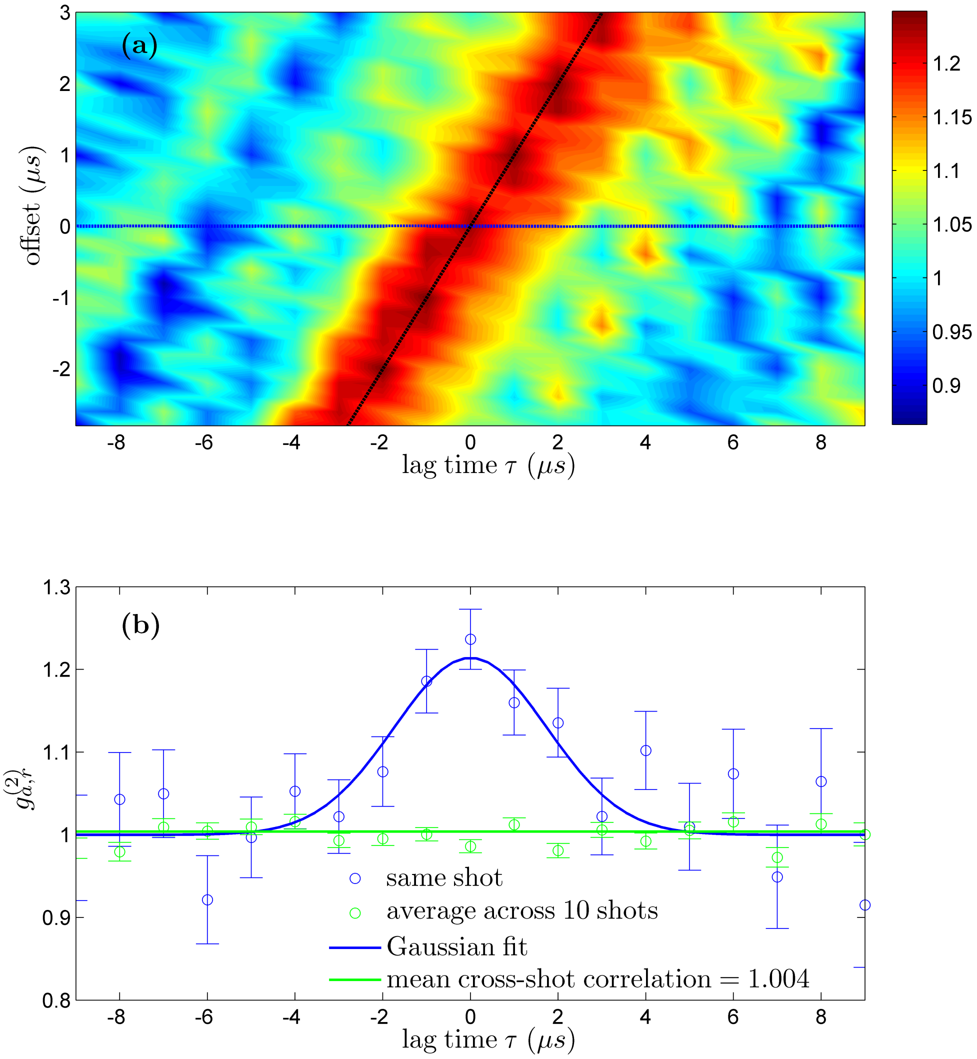

Figure 7.16: The correlation between ASE and RASE temporal modes. (a) $g^{(2)}_{a,r}$ as a function of lag time τ, and also a offset in either the ASE or RASE signal regions - effectively shifting the focal point of the rephasing. The correlation peak is shifted by the same amount and in the same direction as the offset as expected for an echo process, which is highlighted by the black dashed line along τ=offset. (b) The plot of $g^{(2)}_{a,r}(\tau)$ along the blue dashed line marked in (a). The data is fitted with a gaussian peak of amplitude 1.21 ± 0.03 and width $t_c = 4.4 \pm 0.8$ μs. The green points represent the cross-correlation average across 10 adjacent shots, as described in §E.2. This averages to 1.004, meaning that there is very little classical-correlation between shots.

and the signal-to-noise ratios determined from the count histograms. This is higher than the measured value of 1.21, suggesting that the coupling efficiency may have dropped over the course of the measurement.

7.7.2 Auto-correlations

The second-order auto-correlation function is used to quantify the statistics of the individual ASE and RASE emission modes. For low gain, these fields are incoherent and therefore should exhibit a peak in the autocorrelation function, or photon 'bunching' (see §D.3).

Figure 7.17 shows a plot of $g^{(2)}$ as a function of the lag time τ. As expected, the cross-shot correlation average is very close to unity (1.00 ± 0.06 and 1.004 ± 0.009 for RASE and ASE respectively, shown as green and black lines), and so the same-shot $g^{(2)}$ functions are normalized to the overall mean count values as per equation E.7.

The auto-correlation value at $\tau = 0$ is determined with just a single APD (using the algorithm described in Appendix §E.1), however these points are anomalous. In the case of the RASE mode, the value of $g^{(2)}(\tau = 0)$ is much higher than would be expected from the neighbouring points. This is attributed to ringing of the electrical pulse in the connection between the APD and the FPGA board being used to timestamp the counts. This is characterized and discussed in Appendix §F.4. It was found that each detection channel had a higher than expected probability of registering a second count within the 200 ns following an initial count (see figure F.4).

Detector afterpulsing also contributes to false correlations, however these events will occur immediately after the detector dead-time has elapsed, and therefore should not be evident after ~40 ns. As mentioned in §7.6, any counts within 80 ns following a prior count are ignored, to eliminate the effect of afterpulsing, and some of the ringing. However the count rates in the RASE mode are so small, that the remaining spurious coincidences in this $\tau = 0$ time-bin is enough to significantly inflate the value of $g^{(2)}(\tau = 0)$.

In the ASE mode, the auto-correlation value at $\tau = 0$ is actually lower than the level expected (blue points in figure 7.17). A similar ringing effect is occurring in this detection channel, however the count rates in this mode are much larger than for the RASE, and the majority of the detected counts correspond to real photon events. In this case, the post-processing procedure of imposing a 80 ns effective deadtime also eliminates a large fraction of real coincidence events, so $g^{(2)}(0)$ is reduced accordingly. Note that the difference in count rates during the ASE and RASE periods is also reflected in the magnitude of the error bars on the data points in figure E.1; the ASE uncertainty is much less than the RASE because of the larger number of raw counts.

As with the cross-correlations, the data was fitted with a gaussian, however the $\tau = 0$ point was omitted. The RASE peak magnitude according to the fit is 1.66 ± 0.09, and the width represents a coherence time t_c of 4.2 ± 0.5 μs. The ASE peak is slightly lower at 1.46 ± 0.01, and has a coherence time of $t_c = 3.7 \pm 0.1$ μs. These auto-correlation coherence times, and the previously measured cross-correlation coherence time of $t_c =$

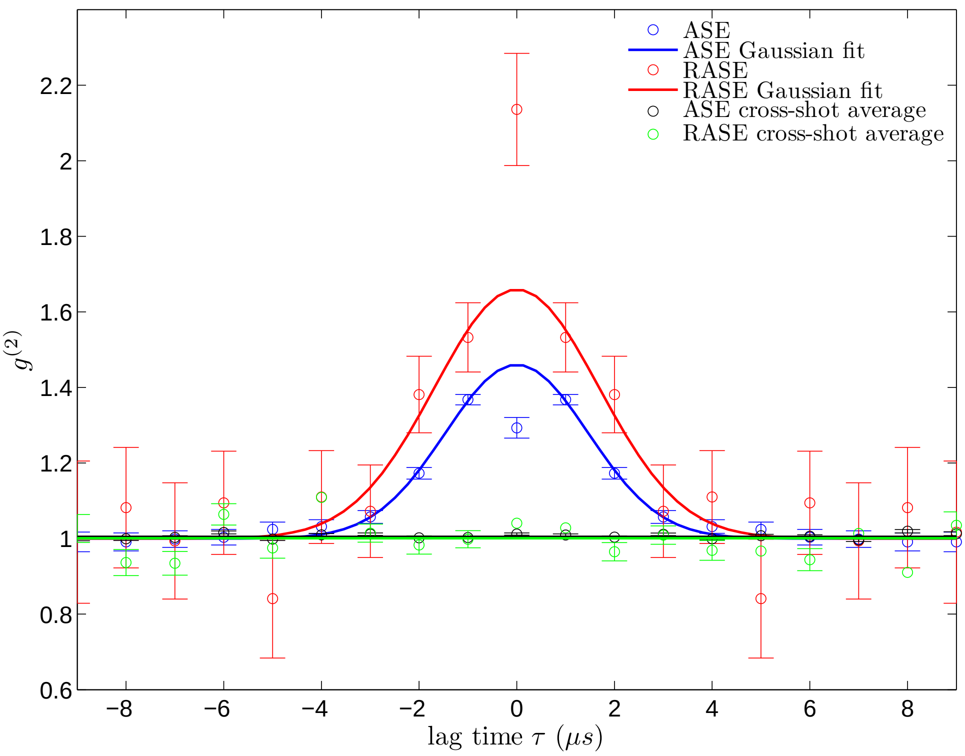

Figure 7.17: Auto-correlation functions calculated for ASE and RASE modes, calculated using a single detector per mode. The point at $\tau = 0$ is artificially inflated for the RASE due to electrical ringing in this detector channel, and is artifically reduced for the ASE due to the post-processing method of ignoring any counts in the 80 ns following a prior count. The ASE count rates are much higher than the RASE, and this post-processing procedure actually discounts many real coincidence events. The gaussian fits to the data are calculated with the point at zero lag omitted. Also shown (green and black) is the correlation between identical spatio-temporal modes of ASE and RASE, but across neighbouring shots. For completely independent signals this cross-shot correlation should be 1. The traces shown are averaged across 10 adjacent shots.

4.4 ± 0.8 μs are all roughly equivalent within the uncertainty of the measurement.

The peak amplitude of the auto-correlation in the ideal case would yield 2, since the ASE and RASE are incoherent fields. However the peak magnitude is reduced due to background noise, losses, and mixing of multiple modes. This can be modelled similarly to the cross-correlation using the theory developed in Chapter 4 with one modification: As discussed in §D.3, the effect of combining multiple modes of thermal states has the effect of reducing the bunching, shown by equation D.16. Since the auto-correlation depends not only on the mean number of noise counts $\langle M \rangle$, but also on the variance of the noise counts, then the effective number of noise modes, N_m, is included.

$$g^{(2)}_{\text{auto}} = 1 + \frac{\langle n \rangle^2 + \frac{\langle M \rangle^2}{N_m}}{[\langle n \rangle + \langle M \rangle]^2} \tag{7.4}$$

where $\langle M \rangle$ is again taken to indicate the mean number of noise counts emitted, which experiences to the same transmission/detection efficiency as the signal. Using the mean count numbers from §7.6 and solving for N_m gives 0.63 for the RASE and 1.2 for the ASE modes. The reduced effective number of noise modes contributing to the RASE signal is smaller than for the ASE is because of the frequency-filtering in the RASE spatial mode.

7.8 The classical-quantum boundary

Having shown a correlation between ASE and RASE, the task remains to determine if the correlation can be explained classically, or if it requires a quantum-mechanical interpretation. There are a plethora of metrics for quantifying non-classical behaviour, two of which were considered in Chapter 5. Considering the low count rates and the photon-counting detection (as opposed to continuous-variable measurement via heterodyne) for the experiment performed here, the ratio of second-order correlation functions R (introduced in §5.1) is chosen as the most amenable metric. Explicitly,

$$R = \frac{g^{(2)}_{a,r}}{\sqrt{g^{(2)}_a g^{(2)}_r}} \tag{5.2}$$

The Cauchy-Schwarz inequality requires that for a classical field model $R \leq 1$. For the values determined in the previous sections, we have measured $R = 1.21/\sqrt{(1.66 \times 1.46)} = 0.77$, which clearly lies in the classical realm of correlations. Thus the expected non-classical correlations between ASE and RASE photons are being completely masked due to the added noise.

7.9 Modifications necessary to measure non-classical correlations

To conservatively estimate the modifications of experimental parameters that would enable measurement of a non-classical correlation, the assumption is made that the auto-correlation functions ($g_a^{(2)}$, $g_r^{(2)}$) take their maximum values. This maximizes the denominator in equation 5.2, and then the required value of the cross-correlation $g_{a,r}^{(2)}$ that would definitively correspond to a non-classical correlation can be determined.

At zero time lag the ASE and RASE modes should exhibit thermal statistics given the incoherent nature of the emission generation (see §D.3). Although the addition of noise will reduce the bunching (as shown by equation 7.4), assuming that the auto-correlations $g_a^{(2)} = g_r^{(2)} = 2$ provides us with a stringent requirement for irrefutably violating this Cauchy-Schwarz type inequality for values of the cross-correlation:

$$g_{a,r}^{(2)} > 2$$

The obtained value of $g_{a,r}^{(2)} = 1.21$ for the experiment described here was well below this mark. By applying the model developed in Chapter 5, the modification required to measure a value of $g_{a,r}^{(2)} > 2$ can be estimated.

Figure 7.18 shows the peak value of $g_{a,r}^{(2)}$ calculated (using equation 7.3) for three cases; first where the both the signal and noise counts are reduced maintaining a fixed signal-to-noise ratio, secondly where just the noise is reduced, and finally where only the signal is reduced. This suggests that reducing the signal and noise counts both by a factor of 10 would enable a measurement of $g_{a,r}^{(2)} > 2$.

However, simply reducing the count rates to reach a regime where non-classical correlations could be observed is not feasible in the current experiment setup. Specifically, the data used to generate figures 7.14, 7.16 and 7.17 represents a total of 9×10^4 shots, which was acquired over approximately 5 hours. To reduce the counts by a factor of 10 and still determine second-order correlations with a similar precision would require increasing the acquisition time by a factor of $\sqrt{10} \approx 3.16$. To maintain consistent optical powers and overlap of all five of the spatial modes for this duration is impractical.

A further problem becomes apparent at low count rates. A second data acquisition of 5×10^4 shots with a reduced gain was performed, and it was found that the noise counts were not reduced by the same ratio as the signal counts. In the analysis of the high-gain data set, emission from the sample was the dominant noise source, and other noise sources are negligible by comparison, however when the signal is reduced this assumption is no longer valid. To be specific, the signal level was reduced by a factor of 5 for the low-gain ASE relative to the high-gain data set. The background noise level during the RASE period was reduced only by a factor of 2.5. Other than detector dark counts, one noise source that is independent of the sample itself is laser light leakage through the fibre of path b. This burns a spectral hole and is detected by the RASE detector.

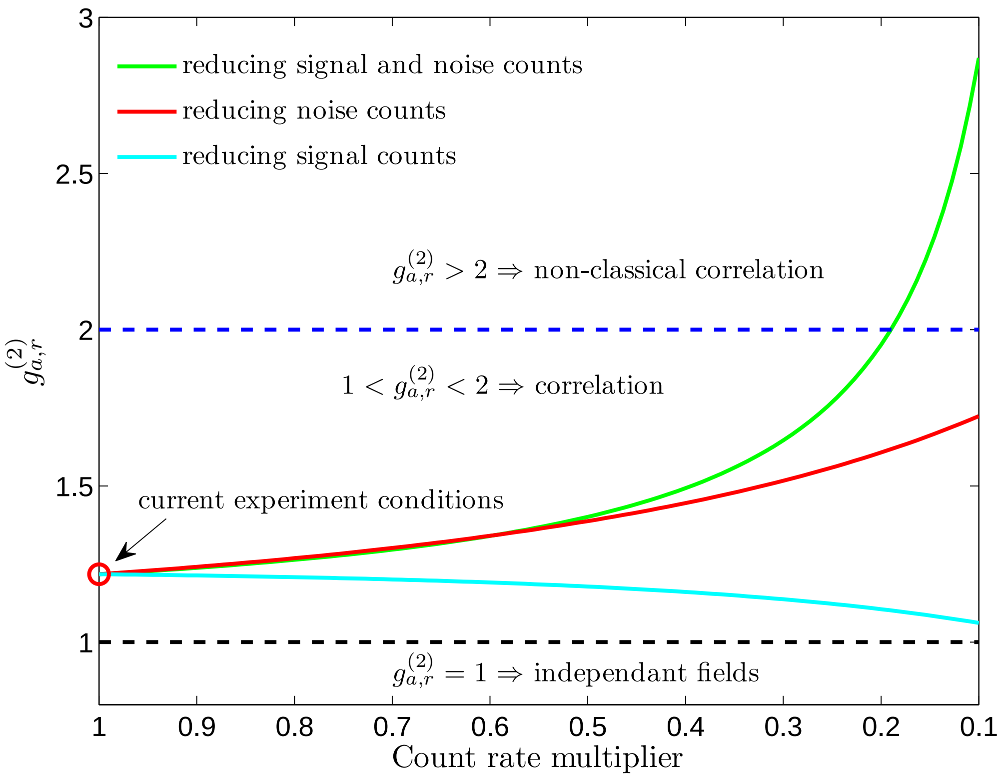

Figure 7.18: Expected trends for the cross-correlation between ASE and RASE modes as the signal and noise counts are reduced calculated using equation 7.3. The green trace represents a decrease of both signal and noise counts, keeping the ratio fixed, whereas the red and cyan traces show a reduction of only the noise (in both ASE and RASE modes) or the signal respectively.

Any reduction of the signal count rate must necessarily be accompanied by a reduction in noise to get closer to the quantum-regime (as shown in figure 7.18). Therefore, to enable a measurement of $g^{(2)}_{a,r} > 2$ in the current experimental setup, simply reducing the number of ions initially inverted will not suffice. In conjunction with reducing the initial population inversion, the modifications that could be made include:

- Frequency filtering of the ASE using another hole-burning filter. (This would require an additional beam-path with frequency control.)
- Reduction of the time taken to prepare the ensemble and filter, which is currently ≈100ms. (This requires more laser power, or a more efficient pumping scheme.)
- Improved transmission efficiencies (for instance by anti-reflection coating the cryostat windows and crystal faces) and rephasing efficiency (by improving mode-matching of the two π-pulse beams).
- Insertion of an additional mechanical shutter to achieve full extinction of the light leakage along path b.
- Improved hole-burnt frequency-filter on the RASE (by increasing the optical-depth), and improved extinction of the red emission using a better interference filter.

7.9.1 Cavity-enhanced RASE proposal

However, even with perfect experimental parameters (ideal transmission/rephasing efficiencies and frequency filtering), there is a fundamental limit on the maximum signal-to-noise ratio achievable. This is determined by the branching ratio of the excited to ground state transitions. As discussed in §7.4, any ions which end up in the ground state $|3\rangle$ *indirectly* via intermediate crystal-field levels can contribute noise in the spatio-spectral-temporal mode of the RASE. Evidence of this signal is seen in figure 7.14. Assuming ideal frequency filtering of all other emission, in the current experiment setup this imposes a maximum achievable value of $g^{(2)}_{a,r} = 3.2$. This would indeed indicate a non-classical correlation, but is somewhat minor in comparison to the degree of non-classicality reported in DLCZ type experiments. For instance in the work of Laurat *et al*, the maximum value of the cross-correlation function was measured as 600 ± 100 [70].

The low quantum efficiency of these long lived transitions is a fundamental issue with rare earth ion systems. The challenge is to somehow increase the desired collective emission (into a given spatial mode) relative to the spontaneous emission background. One way to achieve this would be to modify the geometry of the ensemble. Rearranging the same number of atoms over a larger spatial volume will reduce the solid angle over which there is constructive interference (as described by the function $\Gamma(\phi)$ in §3.3). The signal intensity remains constant but becomes more tightly focussed. Reducing the solid angle

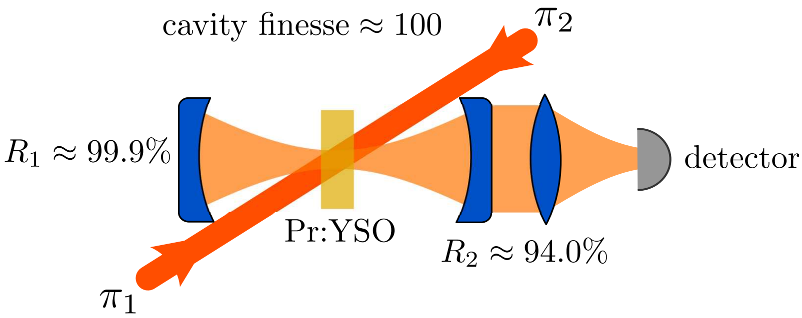

Figure 7.19: Placing the sample for generating RASE inside a Fabry-Perot cavity will improve the ratio of incoherent to coherent population that ends up in ground state $|3\rangle$ by a factor of the cavity finesse. A 100-fold reduction in the noise resulting from this incoherent population should allow for a measurement of a significantly non-classical correlation.

for collection reduces the incoherent background emission, which is radiated in the dipole mode.[17]

Another approach is to modify the transition branching ratio for a particular spatial mode, by placing the Pr:YSO sample inside an optical cavity resonant with the $^3H_4 \rightarrow ^1 D_2$ transition. For example, say that there are N ions in the sample that are initially excited to give a gain G on the ASE mode. If a cavity of finesse f is placed around this sample, as depicted in figure 7.19, then the same gain on the ASE can be achieved with a *reduced* inverted population of N/f. The number of 'incoherent' ions that end up in $|3\rangle$ without emitting an ASE photon is directly proportional to the inverted population, and is unaffected by the presence of the cavity.[18] Therefore the inclusion of a cavity will preferentially enhance the direct transitions at 606 nm, and increase the ratio of coherent to incoherent population that ends up in ground state $|3\rangle$. A cavity of finesse $f = 100$ should be ample.

7.10 Conclusion

In this chapter, we have experimentally verified that spontaneous emission from an ensemble of atoms can be rephased. The experiment was performed with photon counting detection. This poses significant experimental challenges as compared to heterodyne detection due to a lack of spectral discrimination, but is suitable given the ultimate aim of this protocol is to generate entangled photon pairs (and eventually heralded entanglement between remote ensembles).

The major challenge in this experiment was to achieve a high enough signal to noise ratio such that the RASE signal could be discerned. In the first case, this required off-axis rephasing to reduce the FID noise. The four-level echo sequence with a counter-

[17] This can be seen in the equations describing radiation from extended ensembles developed in §4.3.5.

[18] The intermediate crystal-field levels are spectrally broad. This means that population inversion which results in amplification on the kHz-wide transition at 606 nm gives negligible gain on the GHz-wide red emission, and a relatively low cavity finesse of 100 should have no noticeable effect on spontaneous emission in the at the longer wavelengths.

propagating geometry was used to achieve this. This isolated the coherent FID noise, which is approximately 5 orders of magnitude larger than the ASE level, from the detection modes.

The majority of the remaining noise is attributed to spontaneous emission from the sample, but at different frequencies to the RASE signal. To attenuate this emission in the spatial mode of the RASE, further frequency filtering using hole-burning in the back section of the sample itself was implemented. Overall signal-to-noise ratios of $\sim$1 were achieved in both ASE and RASE modes.

A peak in the second-order cross-correlation between ASE and RASE modes confirms the theoretical prediction that amplified spontaneous emission can indeed be rephased. The measured degree of correlation however was not sufficient to confirm its non-classical nature. The correlation is degraded due to the remaining background noise.

In principle it would be possible to improve the experiment parameters such that a quantum correlation could be observed. However, there is a component of the noise that is indistinguishable from the RASE signal, and this imposes a limit on the correlation that would be possible to measure in this experimental setup of $g^{(2)}_{a,r} < 3.2$. This would suffice to demonstrate the non-classical nature of the ASE and RASE photons. However, in light of future plans to use this protocol as a means for generating remote entanglement, then this limited degree of correlation severely restricts the scalability.

In summary, by using a four-level photon echo sequence in a counter-propagating geometry, rephased amplified spontaneous emission has been demonstrated with photon-counting detection methods. Although any quantum-entanglement between the photon pairs could not be verified, through this initial demonstration an understanding of tolerance to experimental imperfections has been gained. It is proposed that performing the RASE experiment in a cavity will negate the noise source which is problematic in the current free-space arrangement.

CHAPTER 8

Dynamic narrow-bandpass frequency filter

This chapter presents a demonstration of a narrow-band frequency filter, which can be switched within µs from being transmissive to absorptive at some particular frequency values. This is achieved through Stark-shifting of spectral holes. This concept was proposed by M. J. Sellars.

The demonstration described here was developed in the context of future RASE experiments. However, this filtering method is a powerful and versatile technique that may be applicable in other applications - most likely for other rare-earth ion based information processing experiments. The methodology developed here represents the only means (to the best of the author's knowledge) to achieve narrow-band filtering (<MHz), with tens of dB extinction, that can be switched within µs.

8.1 Filtering requirements for RASE

As shown in the previous chapter, spectral hole-burning is a useful technique for absorptively filtering out unwanted optical frequencies separated by just a few MHz from the signal. Creating spectral holes however is a slow process, performed by shuffling the population amongst the hyperfine ground states via optically pumping the ions into their excited state, waiting for them to decay back to a ground state, and then repeating the process many times until there is no population remaining in any resonant ground state. Filling in these spectral holes is similarly a process of incoherent population transfer.

In protocols for generating entangled photon pairs, including RASE and DLCZ, the two photons can potentially be generated in the same spatial mode, while being in different spectral and temporal modes. In such geometries, it would be required to have a filter that can rapidly change its spectral properties to transmit each photon in the pair, which are temporally separated by tens of μs, while absorbing at all the noise frequencies.

The techniques for dynamic frequency filtering were developed in the context of the proposed cavity-enhanced RASE experiment, introduced in §7.9. The sample used to generate the RASE is placed inside a Fabry-Perot cavity, and the four-level echo rephasing π-pulses could be applied in one of two geometries. In the co-propagating geometry, the rephasing pulses are input into the cavity mode itself ($\mathbf{k}_{\text{RASE}} = \mathbf{k}_{\text{ASE}} = \mathbf{k}_{\pi_1} = \mathbf{k}_{\pi_2}$), or

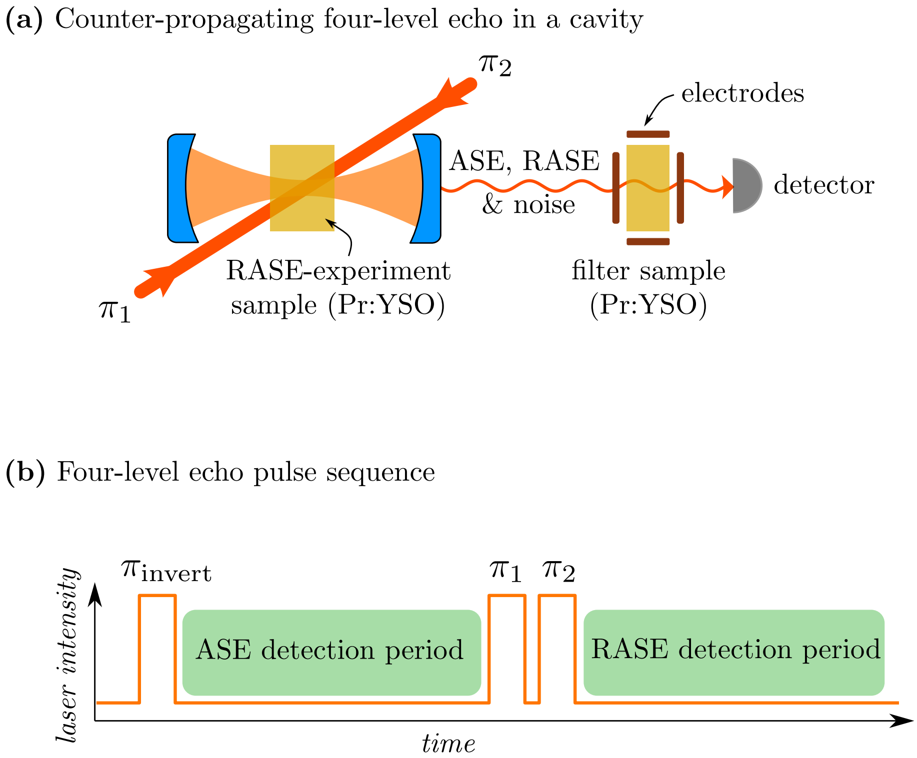

Figure 8.1: A proposed application for the dynamic frequency filter. (a) Outline of a potential configuration of a cavity-enhanced RASE experiment which would require a dynamic frequency filter because the ASE and RASE signals are in the same spatial mode. The π-pulses are shown in a counter-propagating arrangement, however they could also be applied in the cavity mode in a co-propagating geometry. A second sample of Pr:YSO would be used to provide the filtering, with nearby electrodes required to switch the filter properties. (b) shows an outline of the temporal pulse sequence, highlighting the regions where the ASE and RASE signals are emitted.

alternatively the rephasing pulses could counter-propagate ($\mathbf{k}_{\pi_1} = -\mathbf{k}_{\pi_2}$, and $\mathbf{k}_{\text{RASE}} = -\mathbf{k}_{\text{ASE}}$). Either way, the ASE and RASE photons will end up in the same spatial mode as defined by the cavity. Figure 8.1(a) depicts a counter-propagating arrangement.

The filtering would be performed by a second sample outside the cavity, that is programmed to transmit at the ASE frequency during the period following the ensemble inversion, then after the two rephasing π-pulses of the four-level echo sequence, the filter is switched to transmit only the RASE signal.

In a setup where the π-pulses are applied in the same spatial mode as the signal fields (co-propagating), The filtering requirements in the RASE temporal region are particularly stringent since the signal is buried amongst the coherent FID emission which is just tens of MHz away in frequency, and is many orders of magnitude larger in intensity than the signal (see §6.6). This noise source can be avoided using the counter-propagating geometry, where the coherent FID noise is in a different spatial mode to the signal, however this requires a more complicated beam geometry, adding significant experimental overhead and introducing mode-matching inefficiencies.

The purpose of the demonstration described in this chapter is to determine the feasibility of performing dynamic frequency filtering which would be suitable for future cavity-enhanced RASE experiments, and to ascertain which mode of filtering, spectral or spatial, is preferential for distinguishing the signal from the coherent FID noise.

8.2 Combining spectral hole-burning and Stark shifting

The absorption spectrum of the rare-earth ion ensemble can be rapidly changed using the Stark effect. With an applied electric field, the resonant frequency of each ion in the ensemble is shifted proportionally. Therefore by applying a spatially varying electric field across the ensemble, the absorption spectra can be completely reconfigured. Then a dynamic filter can be created by hole-burning spectral features with a particular electric field applied, switching to a different electric field distribution, and hole-burning an alternate spectral structure. After preparing the desired spectral properties in both electric field configurations, then the resulting spectra can be rapidly switched between, simply by switching the external electric field.

This concept combining hole-burning and the Stark effect can be applied to generate as many frequency-transmission windows as desired. Each can be as narrow as the laser linewidth or as wide as the largest ground-state hyperfine splitting, and distributed arbitrarily across the inhomogeneous linewidth of the optical transition. The available attenuation is determined by the optical depth, the width and separation of the spectral holes, and the magnitude and profile of the Stark shift across the length of the crystal. As an example, in the work of Hedges *et al.* [56], the absorption edge of a created spectral feature varied from 0.1 dB to 140 dB in tens of kHz. The potential applications of this type of filter is limited, however, to situations where the field to be filtered is weak enough not to hole-burn itself. This means that the number of photons must be much less than the number of absorbing ions.

The demonstration of this dynamic filter involves hole-burning two spectral trenches, each while a different electric field is applied. This concept is shown schematically in figure 8.2. The aim is to achieve a regime where there is full transmission at ω_1 and a high level of absorption at ω_2 in one electric field configuration, then the field is switched and any signal at ω_1 is heavily attenuated while fields at ω_2 are fully transmitted.

The next section will describe the experimental setup, and following this the Stark shift of the sample and electrode arrangement is quantified. Then for a specific set of filter parameters, namely two 1 MHz wide transmission windows which are separated by 10.2 MHz, the absorption is quantified at the two frequency values in each field configuration using both heterodyne detection and single-photon counting.

8.3 Experimental realization of a dynamic filter

The sample used for this demonstration was $4 \times 4 \times 20$ mm, 0.005% $Pr^{3+} : Y_2SiO_5$, where the beam propagates along the 20 mm length. This sample is held in place with a Delrin support as depicted in figure 8.3, and the electric field is applied using wire wrapped

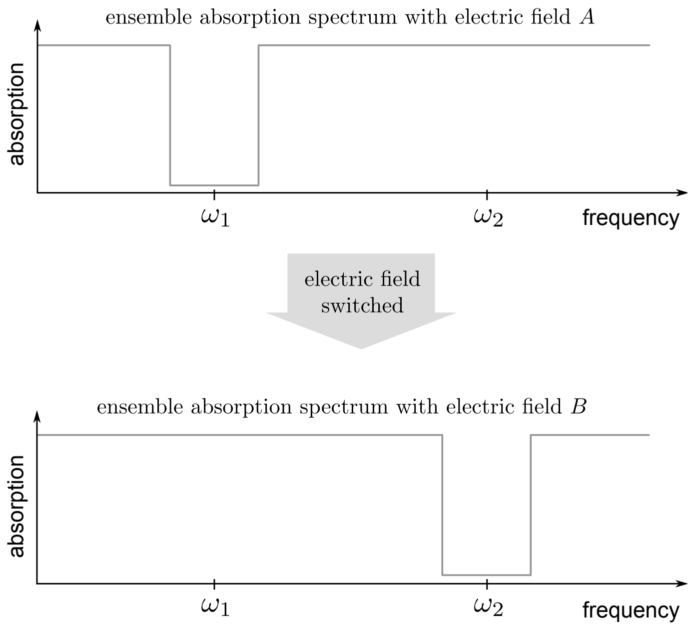

Figure 8.2: Diagram showing the key aspects of the dynamic filtering concept, where reconfiguring the applied electric field switches between two different hole-burnt spectra.

around the sample at either end. One of these wires is grounded (electrode B), while the other (electrode A) is switched up to a maximum of ± 100 V.

This sample holder was designed such that there would be space for a second sample (labelled as the RASE-experiment sample in figure 8.1) in series with this filter-sample in the same cryostat.[1] This alternate sample would be located approximately 2 cm from the filter-sample, in addition to being shielded in a copper Faraday cage, such that switching the voltage of electrode A (at the far end of the filter-sample) would have a negligible Stark effect on the experiment-sample.

The electrode voltage was supplied by an arbitrary waveform generator, the output of which is amplified in two stages by home-built amplifiers to a maximum of ± 100 V, and the maximum switching speed was approximately 3 μs. The electric field as a function of distance along the axis of the sample was numerically modelled using finite-element analysis software (COMSOL). The result is shown in figure 8.3(b). Note that there is no specific requirements for the distribution of the electric field for this demonstration, provided that enough of the ions experience a large enough shift to fill in the spectral holes when the field is switched between the two values.

The different frequency fields that are required for hole-burning in this experiment are all generated using a single double-pass AOM arrangement. A separate path with two AOMs in series is used to generate the local-oscillator field for heterodyne detection. The pulse sequence is shown in figure 8.4. Initially, a broadly sweeping optical field is applied to remove any spectral structure from prior shots. This reset-beam scans over a

[1]While the sample holder used for this demonstration had provision for a second sample, it did not include a Fabry-Perot cavity.

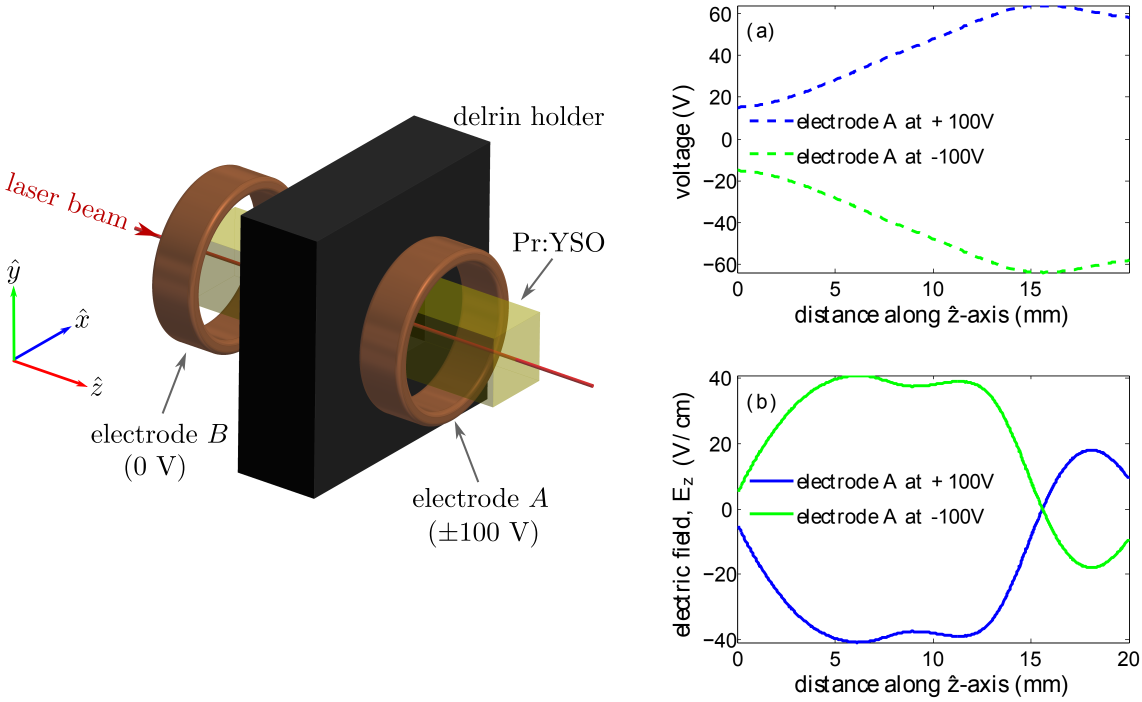

Figure 8.3: Model of the crystal and electrode arrangement used to create a narrow-band dynamic filter. The electrodes were actually coils of wire with a few turns. The sample dimensions are $4 \times 4 \times 20$ mm. Graph (a) shows the voltage along the centre axis of the sample when electrode B is grounded and electrode A is ± 100 V. Graph (b) shows the component of the electric field along the $\hat{z}$ axis, E_z, for the same electrode voltages as in figure (a). Along the beam path, E_y and E_x are zero due to the symmetry of the electrode arrangement. The parameters used for this calculation were; dielectric constant of Delrin $\epsilon_r = 3.7$, for Pr:YSO $\epsilon_r = 10$, and a Stark shift of 110 kHz/(V.cm^{-1}) [91].

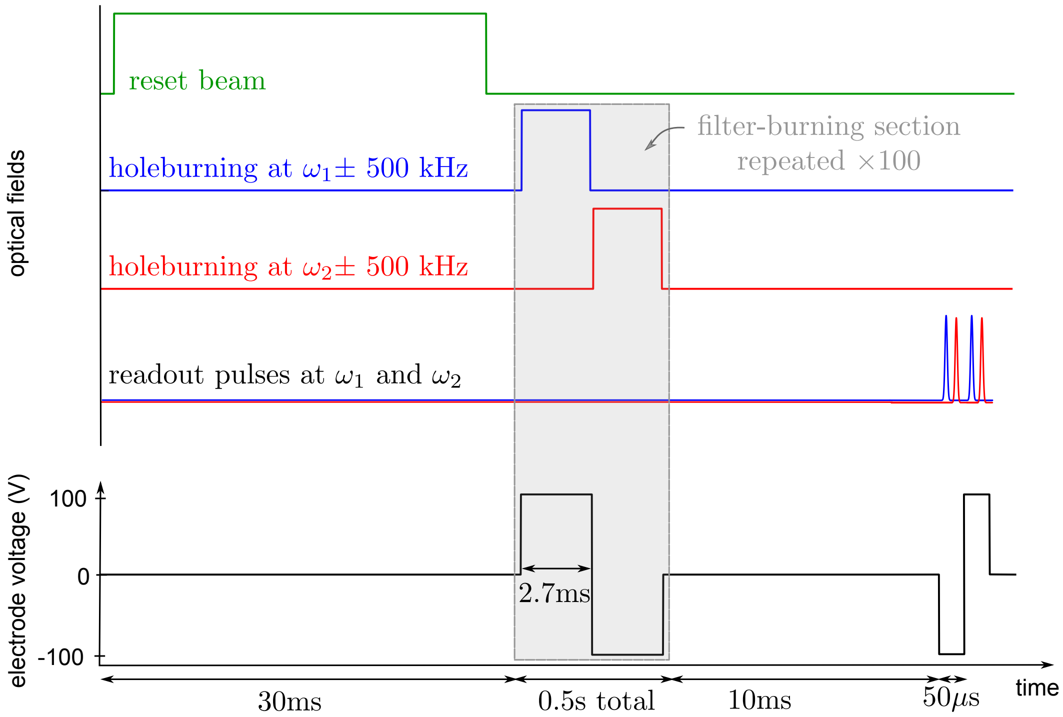

Figure 8.4: The optical field and electrode voltage schematic for preparing a dynamic two-frequency bandpass filter. The reset-beam sweeps over a large section of the spectral profile to erase any residual structure from the previous shot, then the holes at ω_1 and ω_2 are burned while the electrode voltage switches between two values. The filter transmission properties are then read out using temporally short probe pulses at both ω_1 and ω_2 while the electrode voltage is stepped between ± 100 V.

90 MHz range approximately 40 times for a total duration of 30 ms, and both electrodes are grounded during this period.

The filter-burning was performed by switching the voltage of electrode A between +100 V and -100 V every 2.7 ms. When the voltage is high, a field scanning across $\omega_1 \pm 500$ kHz is applied, and when the voltage is switched low, then the beam frequency is scanned over $\omega_2 \pm 500$ kHz, and this process is repeated 100 times. To ensure no cross-talk between hole-burning in different field configurations, the optical fields are off for 50 μs when the electrode voltage is switched. The hole widths of 1 MHz represent a frequency window through which the anticipated signal fields could propagate with negligible dispersion. The frequency difference $\omega_1 - \omega_2$ was chosen to be 10.2 MHz, which is the spacing between the upper two hyperfine levels in the ground state of Pr (see figure 2.1). The entire preparation sequence takes approximately 300 ms.

Following a delay of 10 ms, the filter properties at the two centre frequencies is determined. When using heterodyne detection, the transmission of gaussian probe pulses at ω_1 and ω_2, 900 kHz fwhm in amplitude, are measured for a range of electrode voltage values. The results are confirmed using photon-counting detection.

8.4 Determining the Stark shift

The electric field properties were modelled using finite element analysis (in COMSOL), and the resulting distribution across the length of the sample is shown in figure 8.3(b). Traces are shown for both ±100 V applied to electrode A. To measure the Stark shift obtained with this field profile, a narrow (70 kHz) absorption feature was created in the middle of a large spectral hole (2 MHz).[2] The electric field was then switched on and the absorption profile read out using balanced heterodyne detection of a temporally short probe pulse transmitted through the sample. The results are shown in figure 8.5 for 0 up to 20 V applied to electrode A.

The absorption profile of the initial lorentzian feature which is Stark-shifted as a function of location along the $\hat{z}$-axis can be modelled first by determining the density of ions as a function of the field strength E_z.[3] This is converted to a function of frequency using the Stark-shift in Pr:YSO of 110 kHz/(V.cm), and accounting for the two sub-ensembles of ions which will shift in opposite directions. Then the resulting histogram is convoluted with the initial lorentzian feature profile. The model geometry shown in figure 8.3 predicts an absorption profile which reproduces the main features of the measured profile as shown in 8.5(b). Applying the model to predict the Stark-broadening when 100 V is applied to a group of ions initially all absorbing at the same frequency shows that the population is

[2] The Stark shift could also have been measured using a narrow hole, however using a spectral feature instead gives an immediate picture of how the ions are shifted in frequency, without too much thinking required.

[3] Note that for a quadrupole electrode arrangement as used in ref.[56], E_z is a linear function of distance along $\hat{z}$, and the density of ions as a function of E_z is correspondingly a top-hat distribution. The field profile given by the coil-electrodes used here has a less uniform structure, as shown in figure 8.4 (b).

distributed over a ±4.8 MHz region around the initial absorption feature.

This magnitude of Stark-shifting implies that in switching from -100 V to +100 V the resonant frequencies of the ions will be spread from zero out to ∼10 MHz relative to their starting value. In terms of the filter, where the aim is to fill in spectral holes, this shift should suffice to re-populate the two 1 MHz wide trenches when the field is switched, and thus provide a reasonable level of absorption.

8.5 Characterization using heterodyne

The important properties of the filter to be quantified are the transmission at the two frequencies ω_1 and ω_2 at both electrode voltage values of ±100 V. The amplitude spectra are shown in figure 8.6(b). The dashed trace shows the two probe pulses fully transmitted, recorded with the laser frequency locked 50 GHz from the centre of the $^3H_4 \rightarrow {}^1D_2$ line (which has a linewidth of ≈ 5 GHz fwhm). The peak power of the probe pulses was on the order of 100 nW. The green trace was recorded when the electrode voltage was +100 V, and the laser frequency locked close to the centre of the absorption line, such that the probe pulse at ω_1 (beat frequency with local oscillator of 7.0 MHz) is mostly transmitted, and the probe pulse at ω_2 (beat frequency of 3.2 MHz) is heavily attenuated. The black trace is the same probe pulses for an electrode voltage of -100 V, where the ω_2 pulse is transmitted and the pulse at ω_1 is attenuated.

The level of absorption at both frequencies was also measured at intermediate voltage values to follow the filter transition from one configuration to the other. This trend is shown in figure 8.6(a). Each point represents the average attenuation in dB over a 150 kHz wide frequency window centred at $\omega_{1,2}$. At the +100 V electrode voltage configuration, the pulse at ω_1 is transmitted, with the absorption measured as -0.5 ± 1.7 dB relative to the probe pulse transmission with the laser locked off the line. The pulse at ω_2 is attenuated by 71.2 ± 2.3 dB. Switching to -100V, the pulse at ω_2 is mostly transmitted (5.1 ± 1.8 dB absorption) whereas the pulse at ω_1 is heavily attenuated (67.8 ± 2.2 dB absorption). The difference between the maximum and minimum transmission at each frequency is;

$$68.4 \pm 2.8 \text{ dB at } \omega_1$$
$$66.1 \pm 2.9 \text{ dB at } \omega_2$$

The 5.1 dB of absorption at ω_2 when this pulse should be fully transmitted implies the hole-burning was not sufficient, and should be improved by increasing the length of time spent burning this spectral hole. The difficulty of hole-burning efficiently at this frequency is potentially exacerbated by the fact that the maximum Stark shift when the voltage is switched between ±100 V is roughly equal to the spacing between the two holes, meaning that there are actually some ions that have resonant transitions in *both* the ω_1 and ω_2

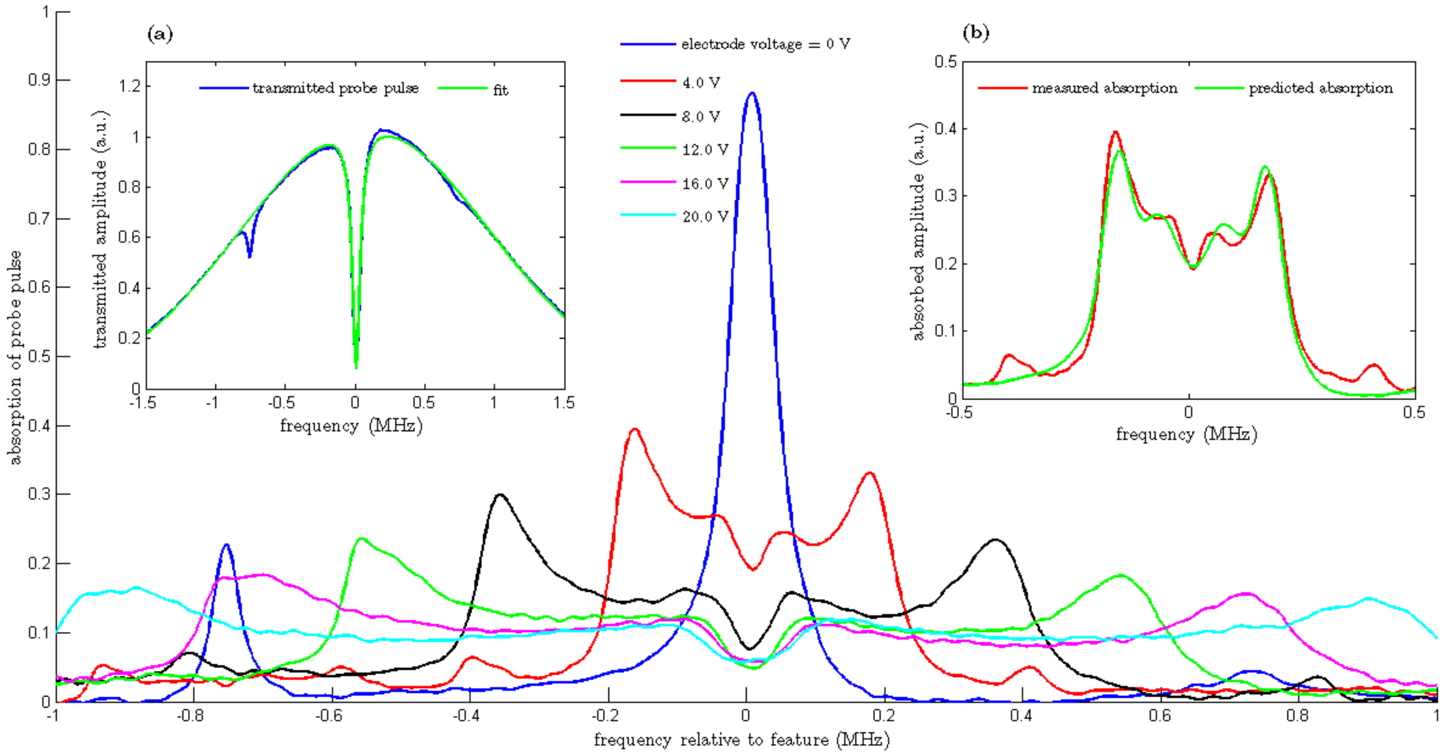

Figure 8.5: Quantifying the Stark shift using a narrow absorptive feature. Inset (a) is the spectrum measured using heterodyne detection of the probe pulse transmission after burning a wide spectral trench, and burning back an absorptive feature. The profile is fit assuming a lorentizan feature, 70 kHz fwhm, subtracted from a gaussian probe profile with fwhm of 2 MHz. The additional absorptive feature located at -800 kHz from the main anti-hole is attributed to some off-resonant burnback of a different frequency class of ions. The main figure shows the absorption profile of the ensemble when a voltage is applied to electrode A (see fig 8.3). The absorption profiles shown are calculated by dividing the measured probe pulse transmission, averaged over 1000 shots for each voltage, by the fitted gaussian pulse input profile. Inset (b) shows the measured and predicted absorption profiles for the electrode A voltage of 4V. The green trace is calculated using the electric field profile shown in figure 8.3(b), and the initial feature profile.

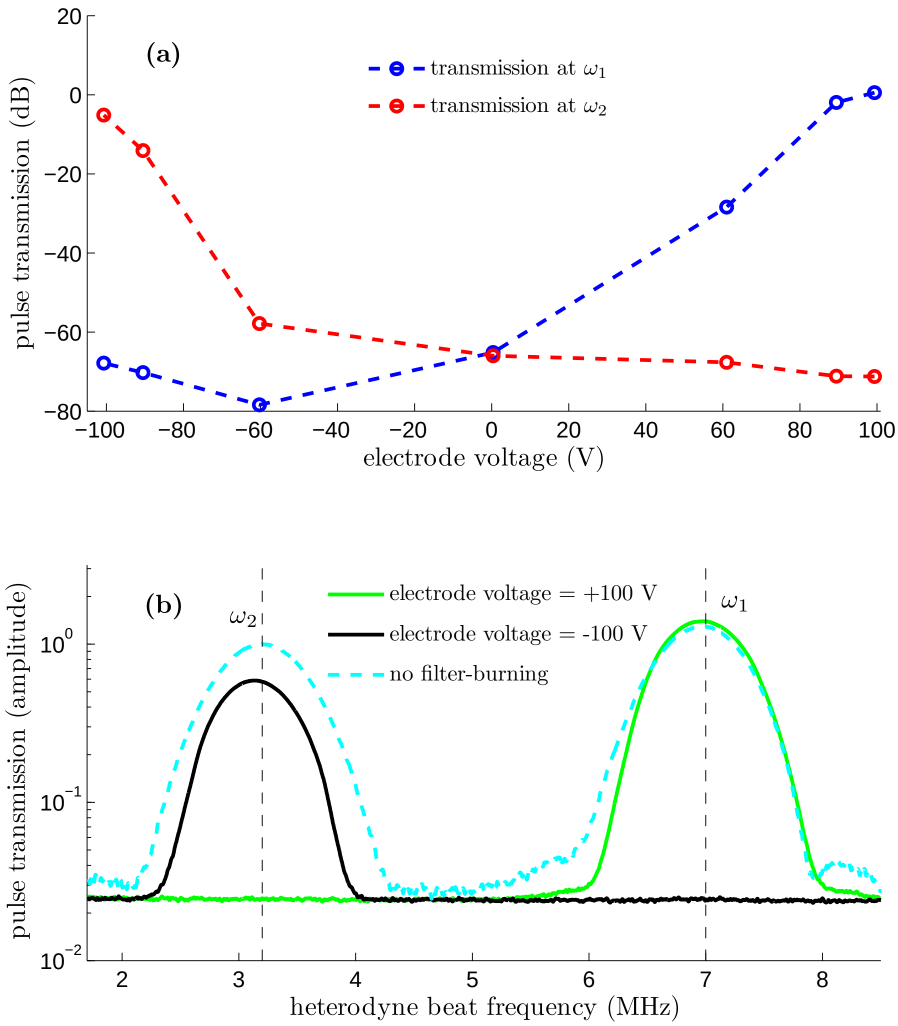

Figure 8.6: Characterization of the filter using heterodyne detection. (a) Attenuation as a function of electrode voltage in a 150 kHz wide spectral region centred at frequencies ω_1 and ω_2. Each point represents the average over 2000 shots, and the error bars are too small to be clearly visible on this scale (maximum uncertainties are 0.7 dB and 0.5 dB for ω_1 and ω_2 traces respectively). (b) The transmission of two probe pulses in the two fully-switched electric field configurations. Also shown is the full transmission of the two probe pulses, obtained by locking the laser 50 GHz from the centre frequency of the $^3H_4 \rightarrow {}^1D_2$ transition.

transmission windows. This is also supported by the fact that the absorption at ω_1 dips to 80 dB with the electrode voltage is just -60 V, and then reduces to 60 dB when the voltage is -100 V (figure 8.6a), indicating that some of the spectral hole at ω_1 is actually Stark-shifted into the ω_2 region when the voltage is fully switched. With sufficient hole-burning it should be possible to remove this particular subset of ions by shelving them in the remaining ground state $|1\rangle$.

8.6 Characterization using photon-counting

Since the purpose of this filter is to provide frequency discrimination, and this is only necessary when the detection method contains no spectral information, the filter performance was also characterized using APDs. In this case, to avoid saturating the detectors, only the single frequency probe pulse which should be heavily attenuated, was input for the ±100 V electrode configurations. Also, the transmission was determined for a range of input pulse powers to check the linearity of the attenuation.

The histogram of counts measured by the APDs is shown in figure 8.7, along with the electrode voltage schematic. The input pulse power at ω_2 is roughly half that at ω_1, and the attenuation was found to be equivalent (within the uncertainty margin) for the three input pulse powers measured. The average attenuation at each frequency was measured as;

$$63.4 \pm 0.6 \text{ dB at } \omega_1$$
$$58.0 \pm 0.7 \text{ dB at } \omega_2$$

This absorption is less than measured using heterodyne, because the transmission measured using APDs is the average over the entire Fourier-width of the probe pulse. The probe pulse fwhm is 900 kHz, however the transmission measured using heterodyne detection was only across the centre 150 kHz region of the holes.

Although there was a mismatch in bandwidth between the heterodyne and photon-counting measurements of the filtering extinction, the difference in absorption at each frequency when the field is switched is surmised to be ~ 60 dB.

8.7 Implications for future RASE experiments

The experiment described in this chapter was a simple demonstration of this dynamic-filtering concept. In applying this technique to the RASE experiment, there are similarly two transmission windows to switch between for the ASE and the RASE emission, however there are multiple frequencies at which attenuation is required (see figure 7.9). This will also be the case for DLCZ experiments.

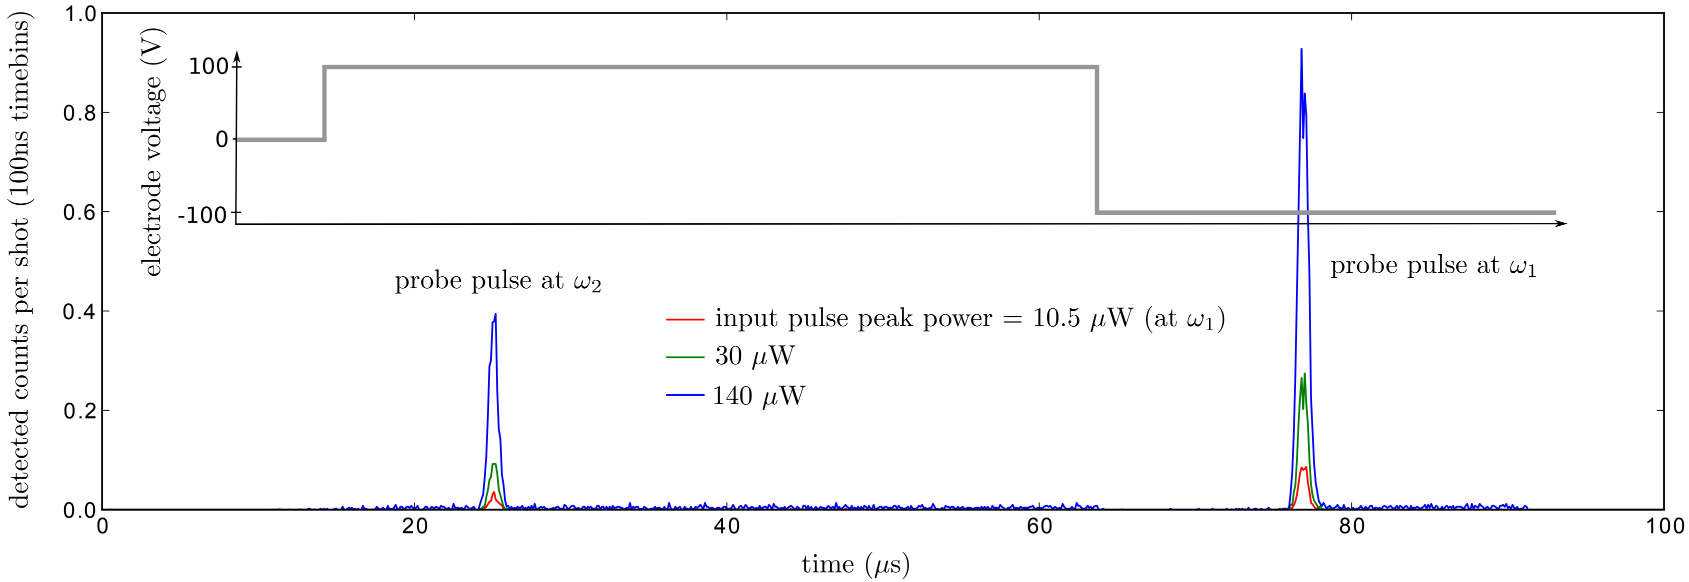

Figure 8.7: Measurement of the filter absorptive properties using an APD. The probe pulse transmission through the dynamic frequency-filter is used to quantify the level of absorption at ω_2 when the electrode voltage is +100 V, and ω_1 with -100 V. Three different probe pulse intensities were used, and are labelled according to the peak power at ω_1. The input power at ω_2 was roughly half that at ω_1. The inset shows the temporal profile of the electrode voltage. The attenuation was found to be ~60 dB, independent of the input power. The total number of shots recorded for the blue, green and red traces were 500, 1600, and 1000 respectively.

In applications with more complicated requirements than the simple two-frequency demonstration described in this chapter, it would be advantages to more carefully engineer the electric field distribution. If the absorption spectra of a Stark-shifted feature shows some structure, as is the case here (see figure 8.5), then there is the risk that a partial spectral hole will be shifted to a frequency where a high absorption is required. One alternative would be to have a quadrupole electrode arrangement such that the electric field component E_z varies linearly along the sample length. The opposite approach would be to have a large constant E_z, to shift all resonant frequencies uniformly, translating a spectral hole to a predetermined frequency region.

In either extreme of a linearly varying or a constant E_z, the larger the Stark-shift the better. A large Stark-shift reduces the overlap between the ensembles of ions that are resonant at both of the transmissive spectral regions. Increasing the field strength however means that it becomes more difficult to shield the nearby experiment-sample from any Stark effects, and may require the filter sample to be moved into a separate cryostat.

The level of absorption measured in this experiment would be sufficient for any RASE experiment where the four-level echo is performed in a counter-propagating geometry, because in this arrangement the source of noise (spontaneous emission to other ground states) is of similar order of magnitude as the signal itself. For example, the experiment described in Chapter 7 included 26 dB of filtering on the RASE mode only, and at this level of extinction this is no longer the predominant noise source.

As discussed in §8.1, this dynamic filter could also be used to block the FID emission if a co-propagating geometry were used for generating RASE. This requires a much higher absorption level compared to filtering of incoherent noise sources, since the intensity of the coherent FID emission will be much larger than the spontaneous emission level. For example the measurements made in Chapter 6 found the FID emission to be 6 orders of magnitude larger than the ASE. In this case the current filter having 60 dB level of attenuation would only achieve a signal to noise of 1. In principle, with the average absorption of 80 dB/cm in the 0.005% Pr:YSO, a higher level of filtering could be achieved by using a different electric field distribution. However the difficulty in doing so (probably requiring locating the filter sample in a separate cryostat) is likely to add more experimental overhead than implementing a counter-propagating geometry for the echo sequence, and spatially filtering the coherent FID emission from the ASE and RASE signals.

8.8 Conclusion

This chapter has presented a simple demonstration of a concept combining spectral hole-burning and Stark-shift properties in an ensemble of praseodymium ions to create a dynamic, narrow-bandpass, rapid roll-off filter. This dynamic filter would be essential for a RASE experiment where the correlated ASE and RASE photons are generated in the same spatial mode. The level of filtering achieved here would be suitable for a cavity-enhanced

RASE experiment with counter-propagating rephasing.

This demonstration involved spectral hole-burning of two transmission windows at ω_1 and ω_2 sequentially, with +100 V and -100 V applied to an electrode near the sample. Switching the electric field distribution reconfigures the absorption spectra of the rare-earth ion ensemble via the Stark shift. Each window was 1 MHz wide, and the separation between the spectral holes was 10.2 MHz. The switching speed was within 3 μs, limited only by the bandwidth of the voltage amplifier. The magnitude of the applied electric field shifts the transition frequency of an individual ion somewhere between 0 and $\pm$10 MHz.

The absorptive properties of this filter were quantified by using heterodyne detection to measure the transmission of a probe pulse at both ω_1 and ω_2 for a range of electrode voltages. The high level of attenuation was also confirmed using APDs, and was found to be on the order of 60 dB. This filter will be useful for RASE- and DLCZ-type experiments for improving the fidelity of generating entangled pairs of photons.

CHAPTER 9

Extending the coherence time of ground-state transitions

This chapter presents a brief discussion of some methods that could be incorporated in the RASE protocol to extend the storage time between the ASE and RASE photons. This discussion is largely based on the experimental work reported in Paper I [1], which has been included as Appendix A. While the paper focuses on the storage of an arbitrary state, in this chapter we consider the implications for specific states relevant for the RASE process.

9.1 Introduction

For interfacing atoms with light, it is necessary to use optical transitions in atomic systems. However, to maximize longevity of a superposition state, transitions between metastable states (with frequencies typically in the RF range) tend to provide longer storage times.

In this chapter we discuss some methods that can be used to extend the coherence time of the metastable transitions, in the context of the RASE experiment. Extending the intermediate storage time between the ASE and RASE photon events is advantageous for the potential quantum repeater applications of this protocol.

This transfer between optical and RF transitions occurs naturally in the four-level echo (4LE) sequence (see Chapter 6). The optical coherence is transferred to the ground states by the first rephasing π-pulse, and the second rephasing π-pulse maps the coherence back to an optical transition.

Once the collective atomic state in the RASE experiment has been transferred to the hyperfine states, various methods can be applied to extend the storage time. The first method considered here is the 'critical point' technique (introduced in §2.4.3). This is a passive method of extending the ground state coherence time, T_2. The second means that will be discussed is an active technique referred to as 'dynamic decoherence control', or DDC [126]. The purpose of DDC is to alter the dynamics of the system on a time scale faster than decoherence occurs. This is implemented by applying many π-pulses to rephase the ensemble state many times within the T_2 timescale. These two techniques can be applied simultaneously.

An experimental characterization of the performance of DDC for maintaining an *arbitrary* quantum state of a spin transition in $Pr^{3+} : Y_2SiO_5$ was reported in Paper I, [1]. The experimental results will not be repeated here, and therefore the paper is included as Appendix A. The purpose of this chapter is to consider the applicability of these techniques to the *specific* states of interest for RASE.

Initially in this chapter, a means of visualizing qubit operations will be discussed. Then the critical point and DDC methods will be briefly revised, before discussing how these methods could be applied in the RASE experiment.

9.2 Geometric picture of qubit operations and qubit errors

The echo sequence applied to rephase the ensemble state can be thought of as an operation on the ensemble 'qubit'. A convenient and intuitive means to visualize the effect of the operations performed is to use the Bloch sphere picture.[1] The entire operation process can be visualized as the transformation of a unit Bloch sphere.

For an ideal quantum memory, this transformation is simply the identity, and the Bloch sphere would remain a unit sphere. However the atoms will inevitably interact with their environment, and therefore an arbitrary qubit state will evolve from being pure (Bloch vector of length 1) to being mixed (Bloch vector < 1). A useful way to categorize the 'errors' that occur in quantum systems is described in reference [127] (and is also utilized in Paper I). Errors are considered in three classes: *unitary*, *incoherent*, and *decoherent*. A unitary process is where the Bloch sphere transformation differs from the desired one by a unitary operation. Incoherent errors are a consequence of inhomogeneity in the ensemble. For instance, the dephasing of a coherent state due to inhomogeneity in the transition frequency is classed as an incoherent error. These errors are reversible by the application of a locally unitary operation. This alters the dynamics of the system so as to 'refocus' the ensemble state. For instance, the effect of inhomogeneous broadening is reversed by applying a π-pulse (a spin- or photon-echo sequence). Finally, decoherent errors are a consequence of fluctuations in the local environment of the individual atoms in the ensemble.

In the following sections we will review the Bloch sphere transformations that result from a simple spin-echo sequence, or a DDC pulse sequence. Since these transformations can be quite different, the metric used in Paper I to compare the performance of different operations was the *volume* of the Bloch sphere transformation. The volume decays due to decoherence, and any not-refocused incoherent errors.

[1] The Bloch sphere representation was discussed in §3.1. Also, see §8.3 of reference [20] for a discussion of quantum operations.

9.2.1 The spin echo

The effect of inhomogeneous broadening of a coherent superposition state can be reversed by the application of a single π-pulse to rephase the ensemble state; a spin-echo sequence. This removes the largest source of incoherent errors, and the remaining decay is due to decoherence. States that are in the $X - Y$ plane of the Bloch sphere (i.e. 50:50 superposition states) will exponentially decay with a time constant T_2, while the population relaxation occurs at a rate given by T_1. The spin-echo sequence in terms of a Bloch sphere transformation is shown in figure 9.1. It is assumed that $T_1 \gg T_2$, which is usually the case for spin transitions.

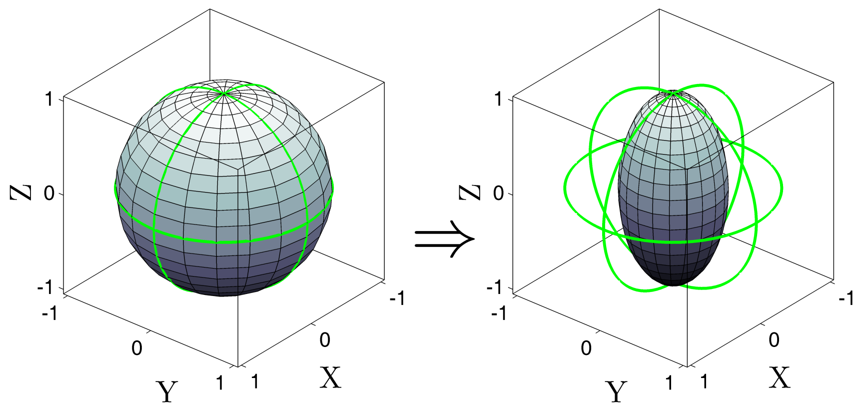

Figure 9.1: Simulated Bloch sphere transformation for a spin-echo sequence. The ensemble state immediately following initialization is pure (unit sphere on the left). Applying a single π-pulse after a delay time τ, rephases the ensemble state after time 2τ (shown on the right). The states in the $X - Y$ plane decohere at a rate given by T_2, while the states at the poles of the Bloch sphere decay at a rate of T_1. For this model, $T_2 = 0.9$ s, $T_1 = \infty$ and $2\tau = 0.6$ s. The volume of the Bloch sphere transformation is 0.26 of the unit sphere volume.

The 'critical point' technique involves applying a specific magnetic field to Zeeman-shift the transition to a turning point. At this point the transition is insensitive to first order to fluctuations in the magnetic field. This maximizes the coherence time T_2.

This technique was originally reported in the work of Fraval *et al.* [89], and applied to the $-\frac{1}{2} \leftrightarrow +\frac{3}{2}$ ground-state hyperfine transition in Pr:YSO. The critical point magnetic field is $\approx$ 780 G for this transition. The same transition was used for the experiments described in Paper I, and the optimal coherence time was measured to be > 900 ms (see figure 4 in Appendix A).

9.2.2 Dynamic decoherence control

The dominant cause of decoherence of the spin transitions in Pr:YSO is due to magnetic interactions with nuclei in the host.[2] The perturbations of the magnetic field at a Pr site occur on a ~10 ms time scale [86]. The concept of DDC is to repeatedly rephase

[2]This was discussed briefly in §2.4. For more detail see the thesis of Fraval [86].

the coherence at a rate faster than the perturbations (and therefore decoherence) occurs. Specifically for Pr:YSO this is enacted by applying rephasing π-pulses less than 10 ms apart.

In the work of Fraval *et al.*, DDC was applied to the $-\frac{1}{2} \leftrightarrow +\frac{3}{2}$ hyperfine transition in $Pr^{3+} : Y_2SiO_5$ and a coherence time of 30 s was reported for a particular input state [128]. For the majority of input states however the coherence was seen to decay rapidly. The reason for this is that the RF π-pulses are slightly inhomogeneous, such that applying many pulses introduces compounding incoherent errors. The Bloch sphere transformation for a DDC sequence is depicted in figure 9.2. This model was generated by assuming a 0.8% linear inhomogeneity in the rephasing pulse area across the ensemble.

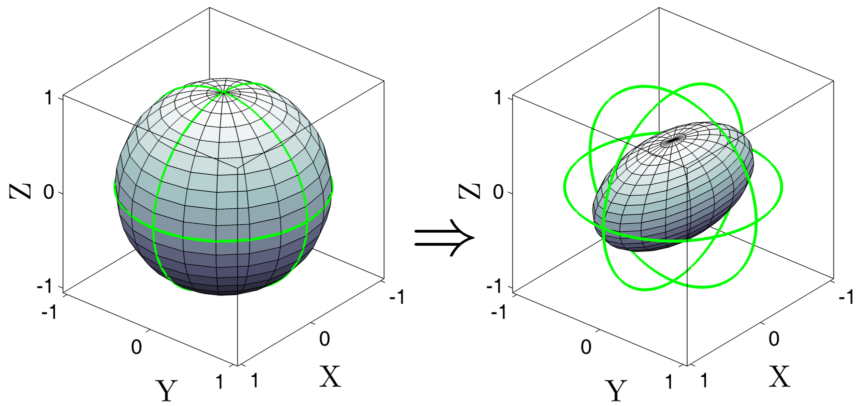

Figure 9.2: Simulated Bloch sphere transformation for a DDC sequence. The pulse sequence is the asymmetric Carr-Purcell sequence depicted in figure 3 of Paper I. This consists of a train of π-pulses, applied every $\tau = 4$ ms. For this model, there is a total of 150 pulses applied, giving a net delay time of 0.6 s (similar to figure 9.1). The inhomogeneity in the rephasing-pulse area for this simulation is a linear distribution with a total variation of 0.8% across the ensemble, and the resulting volume is 0.25 of the original.

Although 30 s is an impressive storage time for a quantum superposition state, the process is limited for general quantum memory applications where the input state is not necessarily known beforehand. The experimental results presented in Paper I were obtained with reduced incoherent errors as compared to the earlier work of Fraval *et al.* [128]. This was a result of three main modifications:

- Reduced inhomogeneity in the RF control field by using a thinner sample. Remaining inhomogeneity was determined to be ∼2%.
- Reduced inhomogeneity in the RF transition frequency, firstly by using a sample of lower dopant concentration, and secondly by Fourier filtering post-measurement. The transition inhomogeneity was approximately 2.5 kHz, however the post-selected 150 Hz component was used to determine the transformation.
- Different rephasing pulses were trialled; specifically designed to be more effective π-pulses for an inhomogeneous control field.

Despite these alterations, it was found that at the critical point field the volume of the Bloch sphere transformation was always degraded by applying the DDC. Applying many rephasing pulses reduces the decoherent errors, however the incoherent errors introduced in doing so were found to reduce the overall fidelity.

9.3 Applicability to RASE storage

The states of interest for the RASE experiment, in the context of a discrete-variable quantum repeater, are Dicke states near the poles of the Bloch sphere. That is, most of the population is in one state, with only a few quanta occupying the second state.

The transformations depicted in figures 9.1 and 9.2 represent similar fidelity (using the volume metric). As a quantum memory for an arbitrary state, these two processes would be considered to perform equally well. However for the RASE experiment, where the specific states to be maintained are near the poles of the Bloch sphere, the shape of the transformations suggest that a simple spin-echo will perform better than DDC.

The compounding errors that are introduced by applying many rephasing pulses with slight inhomogeneity causes mixing of the population between the two levels. In the RASE experiment, when the Dicke state is transferred from the RF transition back to an optical transition, it is imperative that the majority of the population is in the ground state, with only a few quanta associated with the RASE in the excited state. The population mixing caused by performing DDC therefore would add spontaneous-emission noise to the RASE mode.

The issues with implementing DDC effectively are entirely technical in nature. However, the improvement required such that DDC could extend the memory time for RASE would be challenging to achieve. Ideally the control field inhomogeneity should be reduced linearly as the number of applied pulses is increased to ensure there is little mixing of the population states. This would require reducing the level of inhomogeneity measured in Paper I by 2 orders of magnitude. Furthermore, for the RASE experiment the entire inhomogeneous spectral width of the transition should be rephased. To ensure the rephasing pulses are equally as 'hard' for the full 2.5 kHz width as for the 150 Hz width measured in Paper I requires a factor of $\sim$300 increase in the RF power.

Thus, the optimal means of achieving long-term storage in the RASE experiment will likely be to apply the critical point magnetic field, and perform a spin-echo sequence with a minimal number of rephasing pulses. The pulse sequence that achieves this rephasing is shown in figure 6.6. In Pr:YSO this would result in coherence times on the order of one second. Longer coherence times still could be obtained using Eu ions, where the hyperfine ground state transitions potentially have T_2 in the $\sim$100 s range at a critical point field on the order of 2 T [129].

CHAPTER 10

Conclusion and outlook

The main result from this thesis is a confirmation, both experimental and theoretical, that spontaneous emission can be rephased using photon-echo techniques. This was demonstrated in a rare-earth ion doped crystal; $Pr^{3+} : Y_2SiO_5$.

The rephased amplified spontaneous emission (RASE) process was modelled in a fully quantum mechanical framework. This model was simplified by assuming a low gain, however it captures the essence of how and why RASE works in a regime of single-photon detection.

The key features of the RASE protocol can be summarized as follows: Initially the population of an ensemble of atoms is inverted. Following this, the spontaneous emission events are recorded with a single photon counter (in the far field). The measurement process collapses the atomic state into a particular form where the number of atoms in the ground state is known exactly, but there is no information as to which atoms are in the ground state, and therefore the atoms are entangled. This type of state is termed a Dicke state, and can be thought of as a Fock state for atoms.

The Dicke state dephases due to inhomogeneity in the ensemble, but this inhomogeneity can be rephased by the application of a π-pulse, as in a photon echo. The ensemble rephases, and emits a precise number of photons in a well-defined direction. At face value, this may seem somewhat counterintuitive; an initial spontaneous occurrence can be somehow 'reversed' to generate the same precise number of emission events in a well-defined spatial mode. However it was shown that in a fully quantum-mechanical model, the radiative properties of Dicke ensemble states are almost identical (when measuring intensity) to the emission from more familiar coherent atomic states (or Bloch states). This includes, most importantly, the directional nature of the emission, which is simply a consequence of constructive interference of the individual atomic states.

For the experimental demonstration of rephased amplified spontaneous emission (RASE), a novel rephasing pulse sequence was used. This sequence involves two π-pulses applied to a four level atomic system (in a double-Λ configuration), and is termed the four-level echo (4LE). The main motivation for using this rephasing sequence rather than a straightforward two-level echo, is so that the rephasing π-pulses are applied at different frequencies to the single-photon signal fields. As such, the coherent emission following the inevitably imprecise π-pulses (the FID emission) is resolvable in frequency from the signal fields. In addition, the phase matching condition can be satisfied in geometries other than co-propagating, allowing for the FID to also be resolved spatially from the signal.

The experimental investigation of RASE was performed using this 4LE, with a counter-propagating beam geometry. A correlation between the amplified spontaneous emission (ASE) and RASE fields was confirmed. The magnitude of the measured correlation ($g_{a,r}^{(2)} = 1.21 \pm 0.03$) was not enough however to prove its non-classical nature. This was attributed to an insufficient signal to noise ratio. A model was developed to quantify the detrimental effects of background noise and losses on the measured correlation.

Using the counter-propagating 4LE isolated the signal from FID associated noise, and the majority of the remaining background noise in the RASE experiment was spontaneous emission from the sample itself. Most of this emission was at a different frequency to the signal fields, however there was a component of the noise emission that was completely indistinguishable from the signal.

The spectrally-resolvable background noise component was just $\sim$10 MHz from the signal modes. This emission was partially filtered (with $\sim$26 dB absorption) using hole-burning techniques in a different region of the same $Pr^{3+} : Y_2SiO_5$ sample being used for RASE generation. Due to limited resources, this filtering was only applied to the RASE mode, and not the counter-propagating ASE mode. Using spectral hole-burning for frequency filtering is a versatile means of achieving impressive extinction levels.

The component of the noise that is indistinguishable from the RASE signal comes about because of the low branching ratio of the optical transition used for the initial spontaneous emission (the ASE); most of the population that ends up in the ground state ($\sim$97%) is via intermediate atomic levels rather than direct transitions. After implementing the hole-burnt spectral filter, the signal-to-noise of the RASE detection mode was 0.9. The indistinguishable noise component represents roughly one third of the total background noise level detected.

Having identified and quantified the noise sources, suggestions for future experiments were presented that would allow for measurement of a non-classical correlation. First and foremost, the frequency filtering on both ASE and RASE detection modes should be improved. Improved extinction of the background noise that is separated from the signal by $\sim$nm can be achieved with better (higher transmission efficiency and sharper roll-off) bandpass interference filters, which are commercially available.[1] In addition, an integral part of future RASE experiments will be increasing the level of absorption of the holeburnt-filtering.

Modelling suggests that perfect frequency filtering of the signal fields will allow for a non-classical correlation to be measured. Specifically, it is predicted that the best possible correlation between ASE and RASE modes in the current setup would yield $g_{a,r}^{(2)} = 3.2$. However, this is a relatively low 'degree' of correlation, which is fundamentally limited

[1] The filters used in the experiments described in this thesis provided 27 dB extinction, however $\sim$50 dB should be easily attainable with different interference filters.

due to the remaining component of noise that is indistinguishable from the signal.

To negate this noise source, the branching ratio of the transition used for the ASE generation should be increased. One means of achieving this would be to place the sample in a cavity resonant with the ASE transition. The effective branching ratio in the cavity mode would then increase in proportion to the cavity finesse. As such, the requirements for a cavity to achieve a significant improvement in branching ratio are not unreasonable; constructing a Fabry-Perot cavity of finesse 100 for a cryogenic environment is experimentally feasible.

In this proposed cavity-enhanced RASE experiment, the ASE and RASE photons will be co-propagating after out-coupling from the cavity. The detection mode therefore requires a dynamic frequency filter; initially being transmissive at the ASE frequency, and absorptive at all relevant noise frequencies, then switching within µs to being transmissive at the RASE frequency and absorptive elsewhere. In this thesis such a filter was demonstrated. This uses hole-burning to create the transmissive spectral windows, and Stark-shifting to rapidly alter the spectral profile. For the experimental demonstration, the two transmission windows were 1 MHz wide, and separated by 10.2 MHz. In each of these spectral regions, the absorption was switched from $\sim$0.5 dB to 60 dB in a few µs.

Although this dynamic narrow-band frequency filter was developed in the context of the RASE experiment, it is a general technique that could be applied in other rare-earth ion based experiments. The only limitation is that the intensity of the fields to be absorbed must be low enough to ensure they don't cause spectral hole-burning themselves.

The final aspect of the RASE protocol considered in this thesis was the storage time. Long storage times are best achieved using the hyperfine ground state transitions in Pr:YSO. One advantage of the four-level echo sequence is that this intermediate storage on a ground state transition is an inherent part of the rephasing process.

By applying a specific magnetic field, ground-state coherence times on the order of a second have been measured. Previous studies suggest that coherence times can be extended beyond this by using so-called dynamic decoherence control (DDC). However, this applies only for specific coherent states. In general, the application of many control pulses that DDC requires introduces more errors than it removes. Moreover, the incoherent errors introduced are particularly damaging to the specific states of interest for RASE.

To summarize, this thesis presents experimental evidence that spontaneous emission from an ensemble of atoms can be rephased. The non-classical nature of the correlation between the spontaneous emission and rephased fields has been demonstrated by Ledingham in the continuous variable regime [6]. The work presented here has focused on achieving this non-classical correlation regime with single photon counting detection. In this discrete-variable context, a theory was developed utilizing the Dicke state basis for atomic ensembles. This yielded an intuitive description of RASE, connecting the generation of single photons (Fock states) with a method normally associated with coherent states of the field (the photon echo).

In general, rephased amplified spontaneous emission represents a promising and interesting technique for generating entanglement between photons and collective-atomic states.

Appendices

APPENDIX A

Paper on dynamic decoherence control

Demonstration of the reduction of decoherent errors in a solid-state qubit using dynamic decoupling techniques

S. E. Beavan, E. Fraval, and M. J. Sellars
Laser Physics Centre, Research School of Physical Sciences and Engineering, Australian National University, Canberra, Australian Capital Territory 0200, Australia

J. J. Longdell
Jack Dodd Centre, University of Otago, Dunedin 9016, New Zealand
(Received 3 June 2009; published 4 September 2009)

We experimentally demonstrate that decoherent errors for a qubit can be reduced using dynamic decoupling control. The quantum system in our experiment is a praseodymium ground-state hyperfine transition in $Pr^{3+}:Y_2SiO_5$. These experiments were undertaken using an ensemble, where care was taken to reduce incoherent errors resulting from inhomogeneity across the sample. The volume of the Bloch sphere after transformation by the applied pulse sequence was used as a measure of fidelity of the qubit. The strength of the qubit-environment coupling can be tuned by the application of a magnetic field, allowing for regimes of varying decoherence rates to be investigated. We show that the Bloch sphere volume, with the application of dynamic decoupling pulse sequences, decays at a slower rate than observed under free evolution.

DOI: 10.1103/PhysRevA.80.032308 PACS number(s): 03.67.Pp, 76.70.Hb, 42.50.Md, 76.30.Kg

I. INTRODUCTION

Decoherence of quantum systems is a fundamental obstruction in the development of robust quantum information processing technologies. A variety of schemes have been developed, both active and passive, to suppress the loss of quantum information. Active stabilization methods are generally based on quantum feedback and "error-correcting codes" (see, for example, [1]), while passive approaches, or "error-avoiding codes" [2] exploit the symmetry of the system to allow for a subspace of states which will be hardly corrupted, for example, in noiseless-subsystem coding [3].

Dynamic decoherence control (DDC) is an alternative strategy, implemented by applying a time-varying open-loop control which alters the dynamics of the system over time scales faster than the system-environment coupling causes decoherence, effectively "averaging out" the environmental influence [4]. In addition to reducing decoherence directly, an added advantage is that DDC does not require any ancillary measurement or memory resources, and thus can be integrated with other error-avoiding or error-correcting techniques in a straightforward manner to achieve fault-tolerant control.

The theoretical development of DDC is well advanced; however, there are difficulties in applying DDC experimentally. A major constraint on the experimental design is that the control applied to the system must necessarily be faster than the system-environment interaction time.

The nuclear quadrupole transition of praseodymium ions in $Pr^{3+}:Y_2SiO_5$ was recently identified as a suitable quantum system for DDC experiments by Fraval *et al.* [5]. Rare-earth ions exhibit long coherence times in optical as well as hyperfine transitions and, as a dopant in a crystalline host, high optical depths can be obtained. Also, being in a solid state, it circumvents all issues arising from the motion of atoms and increases the available interaction time for measurement or manipulation of the state of the ions.

Previous experiments in a $Pr^{3+}:Y_2SiO_5$ system have demonstrated that decoherent errors could be suppressed to such an extent that coherence times of 30 s were achieved for particular initial states using DDC. However, for the majority of initial states, the coherence was seen to decay rapidly, and this was attributed to inhomogeneity in the ensemble [5]. By improving the homogeneity across the sample, we demonstrate here that DDC offers improvement of fidelity for the qubit as a whole over all input states.

While the experimental investigation of DDC is interesting in its own right, the extended coherence time offered by applying these techniques also adds to the functionality of the rare-earth-doped crystal system for various quantum information processing applications. Experiments in rare-earth-doped solid-state systems to date include stopped light using electromagnetically induced transparency [6] and quantum memory based on switching electric field gradients [7]. There are also implementations of logic gates using the dipole-dipole interaction between ions in the ensemble [8,9] and proposals for incorporating this type of rare-earth memory into the Duan-Lukin-Cirac-Zoller scheme for a quantum repeater [10]. In general, this solid-state system shows promise as a foundation for scalable quantum computing applications, and DDC offers the ability to extend the lifetime of the coherent quantum states.

The process of DDC can be understood by considering the total Hamiltonian describing the evolution of a spin system S and its surrounding environment E [11],

$$H_0 = H_S \otimes \mathcal{I}_E + \mathcal{I}_S \otimes H_E + H_{SE}, \tag{1}$$

where $\mathcal{I}$ is the identity operator. The isolated dynamics of the system and environment are described by H_S and H_E, respectively. The last term accounts for system-environment coupling and gives rise to the decoherence effects. It is the goal of DDC to "switch off" this interaction. The theoretical construction of DDC involves designing a classical timedependant control, $H_c(t)$, which acts on the system alone, such that the effective Hamiltonian describing the total evolution, H_{eff}, no longer contains coupling terms. That is,

$$H_0 \mapsto H_0 + H_c(t) \otimes \mathcal{I}_E \tag{2}$$

such that, for an appropriate (and possibly modified) system Hamiltonian $\tilde{H}_S$,

$$H_{eff} = \tilde{H}_S \otimes \mathcal{I}_E + \mathcal{I}_S \otimes H_E. \tag{3}$$

Physically, the control Hamiltonian applied will not be perfect, particularly since the qubit in this experiment is actually an ensemble of solitary nuclear-spin systems. As such, there will be a variation in the control field seen by each individual ion, in addition to other spatial irregularities. Given this inhomogeneity, a useful way to categorize the myriad of errors that occur in quantum systems is described in Ref. [12]. Errors are considered in three classes: *unitary*, *incoherent*, and *decoherent*. Unitary errors are systematic, and the density matrix describing the system differs from the desired one by a unitary operation. Variation in the system's Hamiltonian H_S and the control Hamiltonian H_c between different ions in the ensemble causes incoherent errors. Incoherent errors (for example, inhomogeneous broadening) are reversible by the application of a *locally* unitary operation to alter the dynamics in such a way that the density matrix is "refocused" to the desired form. Finally, decoherent errors are a result of coupling between an individual spin system and its time-varying environment, which is represented in Eq. (1) through the H_{SE} term. The immediate experimental challenge is to reduce the unitary and incoherent errors to an extent that the effect of the DDC pulse sequences on the decoherent errors can be discerned.

To understand how the incoherent and unitary errors affect the system evolution in isolation from decoherent effects, modeling was performed based on the two-level atom Bloch equations [13]. In this model, unitary errors are a trivial rotation, while the incoherent errors are affected by integrating over a distribution of pulse areas (inhomogeneity in rf field) and detunings (Rabi-frequency inhomogeneity). The decoherent error is represented as a phenomenological decay constant rather than using a full quantum analysis; however, this suffices for this demonstration of DDC; by characterizing and accounting for the unitary and incoherent errors, we can show that the remaining decoherence error in the experiment is reduced by the application of DDC.

II. METHODS

The sample used is a 1-mm-thick 0.005% $Pr^{3+}:Y_2SiO_5$ crystal, cooled to 2 K. There are two inequivalent sites in the Y_2SiO_5 lattice where the Pr^{3+} ions can substitute for Y, and in this work we use *site 1* ions, as labeled in Ref. [14]. As is the case for many of the rare-earth ions, the rate of relaxation between the ground-state hyperfine levels is very slow, on the order of days [15]. Superposition states decay much quicker, predominantly due to environmental interactions. The host Y ions have thermally populated spin states, and cross relaxation occurs between these ions via the dipole-dipole interaction. This causes small but frequent ($\tau_c \sim 10$ ms) fluctuations in the magnetic field at the Pr nuclei, and thus forms a decoherent error source. However, the Y ions in the immediate vicinity of the Pr sites are perturbed by the presence of the large magnetic moment (Pr magnetic moment is 10 kHz/G [16], compared to 100 Hz/G for Y [17]) and no longer resonantly exchange spin with the bulk Y ions. It is reasonable to expect that cross relation between these ions is so slow as to be static on the ms time scale of our experiments. This *frozen core* field varies from ion to ion and thus is characterized as an incoherent error source.

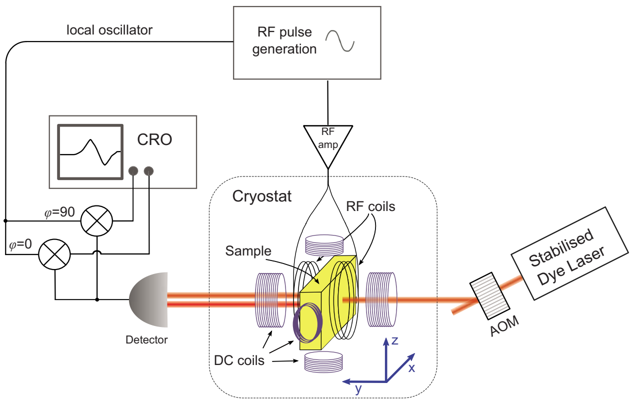

FIG. 1. (Color online) Experiment setup. The $Pr^{3+}:Y_2SiO_5$ sample is in a cryostat along with superconducting coils along x, y, and z directions to provide the critical-point magnetic field, and rf coils to apply the various pulse sequences. Raman heterodyne detection is used to read out the state of the ensemble, where an acousto-optic modulator is used to gate the light. The signal from the photodiode is mixed with a local oscillator to extract the signal at the transition frequency of 8.65 MHz.

An advantageous attribute of the $Pr^{3+}:Y_2SiO_5$ system is that the sensitivity of the ions to environmental magnetic fluctuations can be controlled by applying a static magnetic field [18]. To maximize the coherence time, the applied field is chosen to Zeeman shift the transition frequency such that it has no first-order sensitivity to magnetic field variations (∂ transition frequency$/\partial B_x, B_y, B_z = 0$). This extends the lifetime of superposition states, the coherence time T_2, significantly. In a coordinate system where y is parallel with the c_2 axis, z is the direction of the predominate polarization of the optical $^3H_4 \rightarrow {}^1D_2$ transition, and x is perpendicular to both (Fig. 1); the *critical-point* magnetic field is $B_{CP} = -732$, -173, and 219 G for the $m_I = -\frac{1}{2} \leftrightarrow +\frac{3}{2}$ transition. This is applied using three pairs of superconducting magnets aligned along the x, the y, and the z axes and, at the critical point, T_2 of 900 ms can be achieved. Also, tuning the magnetic field around the critical point offers a means to control the sensitivity of the ions to the environmental perturbations (decoherent errors), and thus we can study a range of dynamical regimes with varying relative contributions of incoherent and decoherent errors.

Another feature of this system is the high signal to noise ratio of the ensemble spin measurements. Although the transition of interest is in the rf frequency range, the qubit initialization and the final-state detection are done optically using a Raman heterodyne technique [19]. This is applicable in

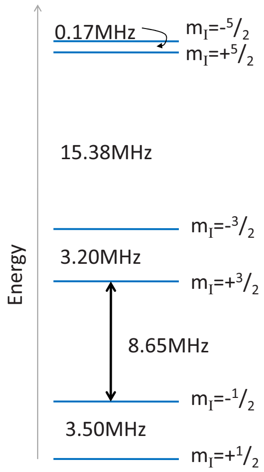

FIG. 2. (Color online) Ground-state energy levels of Pr^{3+}: Y_2SiO_5 at the critical-point field. Energy splittings are shown, and the qubit transition is indicated by the arrow.

systems where the two hyperfine-split ground states have a common excited state (a Λ system). In our case, the transitions to the excited state occur at optical frequencies. Driving one optical transition, in conjunction with existing coherence between the two ground states, will establish a coherent relationship between all three levels and generate the second optical field. The amplitude of the beat signal between the two optical frequencies is directly proportional to the coherence between the ground states.

The laser used was a *Coherent 699* frequency stabilized dye laser, modified to have a subkilohertz linewidth. The wavelength was tuned to the $^3H_4 \rightarrow {}^1D_2$ transition at 605.977 nm. The light incident on the sample had a 100 μm spot diameter, gated using a 100 MHz acousto-optic modulator, and detected with a fast photodiode (Fig. 1). By mixing this signal with a local oscillator referenced to the rf driving field, the Raman heterodyne signal was recovered. Quadrature detection was facilitated by using two phase channels of the local oscillator $90° \pm 0.5°$ apart.

The level of sensitivity offered by this optical detection system allows for very dilute dopant concentrations to the extent that the Pr-Pr interactions are negligible and each ion can be considered as an isolated quantum system. At equivalent separation, the interaction strength between two Pr ions is 50 times larger than that of the Pr-Y interaction; however, the Pr concentration of 0.005%, and hence a large average separation distance, means that the Pr-Pr interaction in the sample is effectively ~400 times smaller than the Pr-Y interaction.

The rf signal to drive the qubit transition (see Fig. 2) was generated by a synthesized function generator and split six ways to provide the local oscillator signal and five phase channels necessary to produce all the pulse sequences for this experiment. Phase delays of 0°, 90°, 104°, 180°, and 313° (accurate to <1%) were achieved using different cable lengths, and appropriate attenuation of each channel was applied to coarsely correct for different intensities. Transistor-transistor logic (TTL) outputs of a direct digital synthesis system were used to pulse the rf and switch between the phase channels. This rf pulse sequence was amplified by a 200 W amplifier, with gain sufficiently reduced to ensure a linear response, and applied to the sample using a six-turn coil with a diameter of 8 mm. Scaling the length of the pulses with a resolution of 10 ns gave fine-tuning control to ensure consistency in the pulse area through the different phase channels. The Rabi frequency was found to be 32 kHz on average, with less than 1% variation between phase channels.

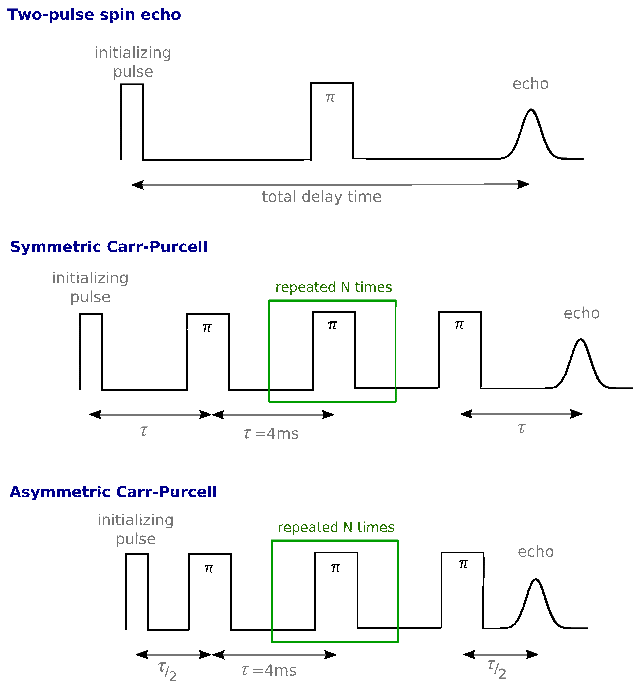

FIG. 3. (Color online) Pulse sequences used in the experiment. The two-pulse spin echo rephases inhomogeneous errors, while the Carr-Purcell sequences that aim to additionally reduce decoherence by rephasing the ensemble faster than environmental perturbations occur.

Prior to applying each pulse sequence, a combined optical and rf repump scheme was used to prepare a pure-state ensemble. While optically radiating the sample, rf frequencies of 12.15, 15.35, 15.55, and 18.59 MHz were sequentially pulsed with a duty cycle of 10% and the entire cycle was repeated ~200 times. These rf frequencies drive transitions from five of the six hyperfine ground states, and in conjunction with the optical field driving transitions to the excited state, after many cycles the ion is likely to relax to the unconnected $m_I = -\frac{1}{2}$ state. The entire burn-back sequence duration is <100 ms.

The pulse sequences used in the experiment are depicted in Fig. 3. The simple two-pulse spin echo was measured as a gauge of the free evolution of the system with incoherent errors removed. This sequence consists of an initial $\frac{\pi}{2}$ pulse to place the sample into a superposition state; then after the system has evolved freely for time τ, a π pulse is applied to *refocus* the Bloch vector components and after further time τ an echo of the original coherence is obtained. The DDC pulse sequences are analogous to the classical *bang bang* control or Carr-Purcell (CP) sequences in NMR. This train of

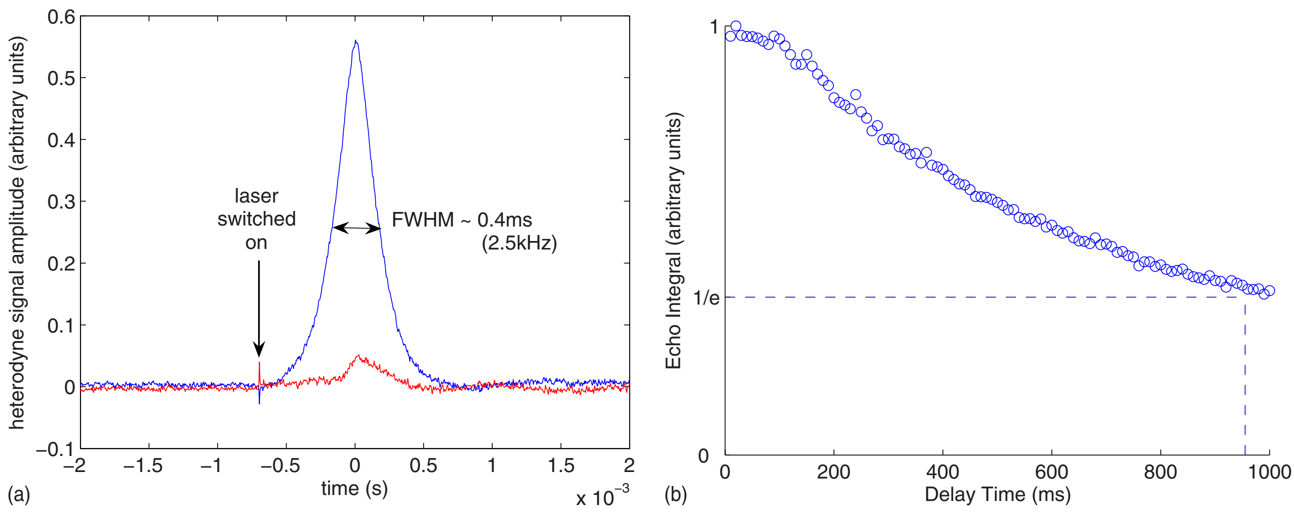

FIG. 4. (Color online) (a) Oscilloscope screenshot of the Raman heterodyne signal. The signal from the photodetector is filtered to select the component at 8.65 MHz. The different traces (red and blue curves) are the 90° separated phase detection channels. (b) Echo integral as the total delay time is increased. Since the static incoherent errors are rephased by this pulse sequence, decay is due to decoherent errors. The time constant characterizing the decay, $T_2=950$ ms, is also marked.

π pulses is intended to rephase the coherence within the time window of the environmental perturbations.

III. SYSTEM CHARACTERIZATION

The frequency of the $m_I=-\frac{1}{2}\leftrightarrow+\frac{3}{2}$ hyperfine transition was found to be 8.6522 MHz, with an inhomogeneous linewidth of 2.5 kHz (Fig. 4). This linewidth was significantly narrower than the 4 kHz observed by Fravel *et al.* [5], attributed to use of a sample with a lower concentration of Pr ions, which is expected to exhibit less strain induced broadening. Second, the thickness of the sample in the direction of propagation was reduced from 3 to 1 mm. This could also result in a narrower line as a result of improved homogeneity in the dc magnetic field across the sample. Additionally, the recorded echo was Fourier filtered postmeasurement to select a subgroup of ions within the inhomogeneous line with a narrower ($\sim$150 Hz) detuning from the central Larmor frequency.

At the critical-point field, where decoherent errors are minimized, the coherence time of the ensemble was found to be $\sim$900 ms, determined using the two-pulse spin echo [see Fig. 4(b)]. The duration of a pulse required to achieve a pulse area of π was 15.6 μs.

A significant incoherent error arises due to technical errors in the experiment, such as inhomogeneity in the applied rf. To characterize this, a spin nutation experiment was performed, again using the two-pulse spin-echo sequence, with an increasing duration of the initial pulse (from 0 to 1.5 ms or 100π) and a constant relatively short delay time (10 ms or 1.1% of T_2). The decay of the nutation is indicative of a variation in the rf intensity across the sample and a corresponding distribution of pulse areas seen by individual ions. The nutation envelope was fitted assuming a quadratic variation in the rf intensity across the sample. It was found that this variation was approximately 1% of the maximum rf intensity (see Fig. 5).

IV. RESULTS

A. On the critical point

A single process tomography measurement for a given input state requires a minimum of nine measurements [20]: three initial states $[X=\frac{1}{\sqrt{2}}(\langle 0|+\langle 1|)$, $Y=\frac{1}{\sqrt{2}}(\langle 0|+i\langle 1|)$, and $Z=\langle 1|]$, measured along three orthogonal axes. For example, measurement of the spin echo of a system initialized in state X after a short delay time would return a normalized result close to $X'=(1,0,0)$ along (X,Y,Z). However, as the delay time is extended, and the influence of unitary, incoherent, and decoherent errors becomes apparent, there are nonzero components of the echo along Y and Z as well as the diminishing X component. A 3×3 transformation matrix for a particular delay time is obtained from the measured vector transformations $X\rightarrow X'$, $Y\rightarrow Y'$, and $Z\rightarrow Z'$. Assuming that any arbitrary initial state $A=(x,y,z)$ transforms according to its components [21] as $A'=(x',y',z')$, this transformation can be visualized as a three-dimensional (3D) Bloch sphere image.

Figure 6 shows how the Bloch sphere is transformed with increasing total length of the pulse sequence for three different rephasing pulse scenarios. These transformations were normalized to the map for the shortest delay time in order to remove artifacts of systematic errors in the detection system.

The first pulse series (column 1 of Fig. 6) is a spin-echo sequence with increasing delay time. The evolution of the Bloch sphere in this case is straightforward; the decay of the coherent states is characterized by T_2, and the population states persist indefinitely since T_1 is much larger than the time scale of this experiment. Since the echo of the ensemble spin is obtained due to the rephasing of incoherent errors, then this series represents decay induced by decoherent errors alone and serves as a standard for comparison to determine the effectiveness of the dynamic techniques.

The first DDC sequence, the Carr-Purcell sequence ($CP_{\pi/2}$), is simply a train of uniform phase π pulses separated

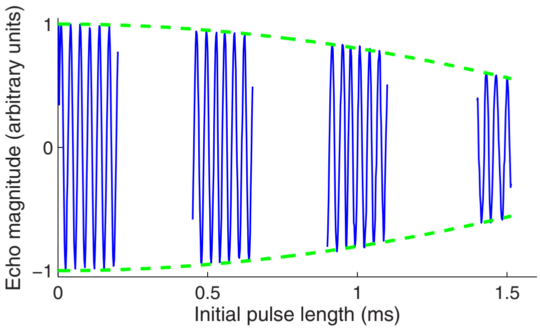

FIG. 5. (Color online) Spin-echo magnitude for nutation sequence; initial pulse of increasing duration (x axis) followed by a π pulse applied after short delay time (10 ms) and echo signal measured after further 10 ms wait. To see significant decay, the measurement was broken into four sections. The damping is attributed to inhomogeneity in the rf intensity. The envelope of the nutation was fitted with the Fourier transform of a step function describing the Rabi frequency of the ions in the sample as a quadratic function of their position.

by cycling time τ (see Fig. 3). With an initial delay time of $\frac{\tau}{2}$, the sample coherence is rephased midway between each π pulse, and is thus referred to as symmetric. An asymmetric variation on this (CP_τ) has an initial delay time equal to the cycling time τ, such that the spin echo coincides with every second π pulse. The cycling time was chosen as 4 ms, which is within the 10 ms time scale of the host ion spin flips, and therefore is in the regime of dynamic decoupling.

The symmetric sequence, seen in column 3 of Fig. 6, preserves Bloch vectors along the Y axis (axis of rotation) almost perfectly, while the fidelity of the X and the Z states is limited by the effect of compounding errors in pulse area over the course of the sequence. This is consistent with the optical Bloch equation model (see Fig. 7), which shows that integrating over a range of pulse areas causes the decay of the X and the Z Bloch vector components and also a slight rotation about the Y axis. The results for the asymmetric sequence (column 2) demonstrate reduced rephasing of initial states along the Y axis and improved rephasing of the X-axis states. This is also in agreement with the model (Fig. 7).

As a qualitative comparison of the pulse sequences, we use two measures of fidelity: the volume and the rotation of the Bloch sphere transformation. Conventional figures of merit used to describe the fidelity of quantum processes are usually a combination of these; for instance, for the single-qubit case the *gate fidelity* as described in Ref. [12] is

$$\text{Gate Fidelity} = \left\langle \frac{\boldsymbol{R}_{ideal}}{|\boldsymbol{R}_{ideal}|} \cdot \boldsymbol{R}_{measured} \right\rangle, \tag{4}$$

where $\boldsymbol{R}$ is a Bloch vector and $\langle\ \rangle$ denotes averaging over all input states. Here, the dot product will scale the fidelity for a rotation error, while the magnitude of $\boldsymbol{R}_{measured}$ will represent the decay of the Bloch vector. In our case, the ideal transformation is simply the identity, and any measured rotation is predominantly due to technical errors in the applied pulse sequence [22]. Therefore, determining the volume and the rotation of the Bloch sphere transformation is an intuitive way of monitoring the decay of the Bloch sphere, caused by decoherent and incoherent error sources, *separately* from the unitary error. Moreover, the volume fidelity reflects the equal importance of conserving all components and will be reduced to zero if a single dimension is lost, which reflects the ineffectiveness of such a system as a quantum memory. This is notably dissimilar from using the average Bloch vector length, which will be nonzero even in the case when only a purely classical memory, a one-dimensional line across the Bloch sphere, remains.

The 3×3 matrix ($\mathbf{M}$) describing the Bloch sphere transformation can be factorized as

$$\mathbf{M} = \mathbf{U} \times \mathbf{A} \times \mathbf{V}^{\mathbf{T}}, \tag{5}$$

where $\mathbf{A}$ is a diagonal matrix and $\mathbf{U}$ and $\mathbf{V}$ are rotation matrices. When applied to the unit sphere, the effect of the matrix $\mathbf{M}$ can be interpreted geometrically as a rotation by $\mathbf{V}^{\mathbf{T}}$, followed by a scaling of the axes by $\mathbf{A}$, and a further rotation $\mathbf{U}$. The volume of the resulting ellipsoid is determined by taking the product along the diagonal of $\mathbf{A}$. A measure of the rotation is obtained using the matrix $\mathbf{U}\times\mathbf{V}^{\mathbf{T}}$, which describes the total rotation with the scaling of the axes removed. An arbitrary 3D rotation, such as that described by $\mathbf{U}\times\mathbf{V}^{\mathbf{T}}$, can be written in terms of an axis of rotation and an angle, and here we will use this angle as a measure of rotation fidelity.

Figure 8 shows the volume and the rotation fidelity trends for the various pulse sequences and clearly shows that the symmetric and the asymmetric CP sequences are reducing the fidelity. Applying the critical-point technique minimizes decoherence to such an extent that the incoherence generated by the application of many pulses becomes the dominant source of error. Therefore, in this regime, the effectiveness of the DDC at reducing the decoherent error is indeterminable.

To quantify the level of technical noise, the model was fitted to the CP volume decay trend by integrating over a ±150 Hz detuning range from the central Larmor frequency and a 2% linear range of pulse areas. Although the nutation decay only suggested a 1% variation in the rf amplitude, the increased error seen here is attributed to stochastic errors introduced by the rf amplifier, the output of which was observed to fluctuate by $\sim$0.7% from pulse to pulse.

B. Away from the critical point

To reach a regime where the decoherence is visible in addition to incoherence, we can simply tune the magnetic field away from the critical point. In this way, the sensitivity of the ions to the environmental quantum perturbations, or equivalently the sensitivity to decoherent errors, is increased. Conversely, the incoherent error remains unperturbed since the main contributing factor to inhomogeneous broadening is crystal strain [23], which is independent of the externally applied magnetic field.

In the first instance, the Z component of the magnetic field was reduced by 0.27% to 218 G. The coherence time is reduced to 380 ± 10 ms; however, other properties of the system are practically unaffected; the frequency of the transition

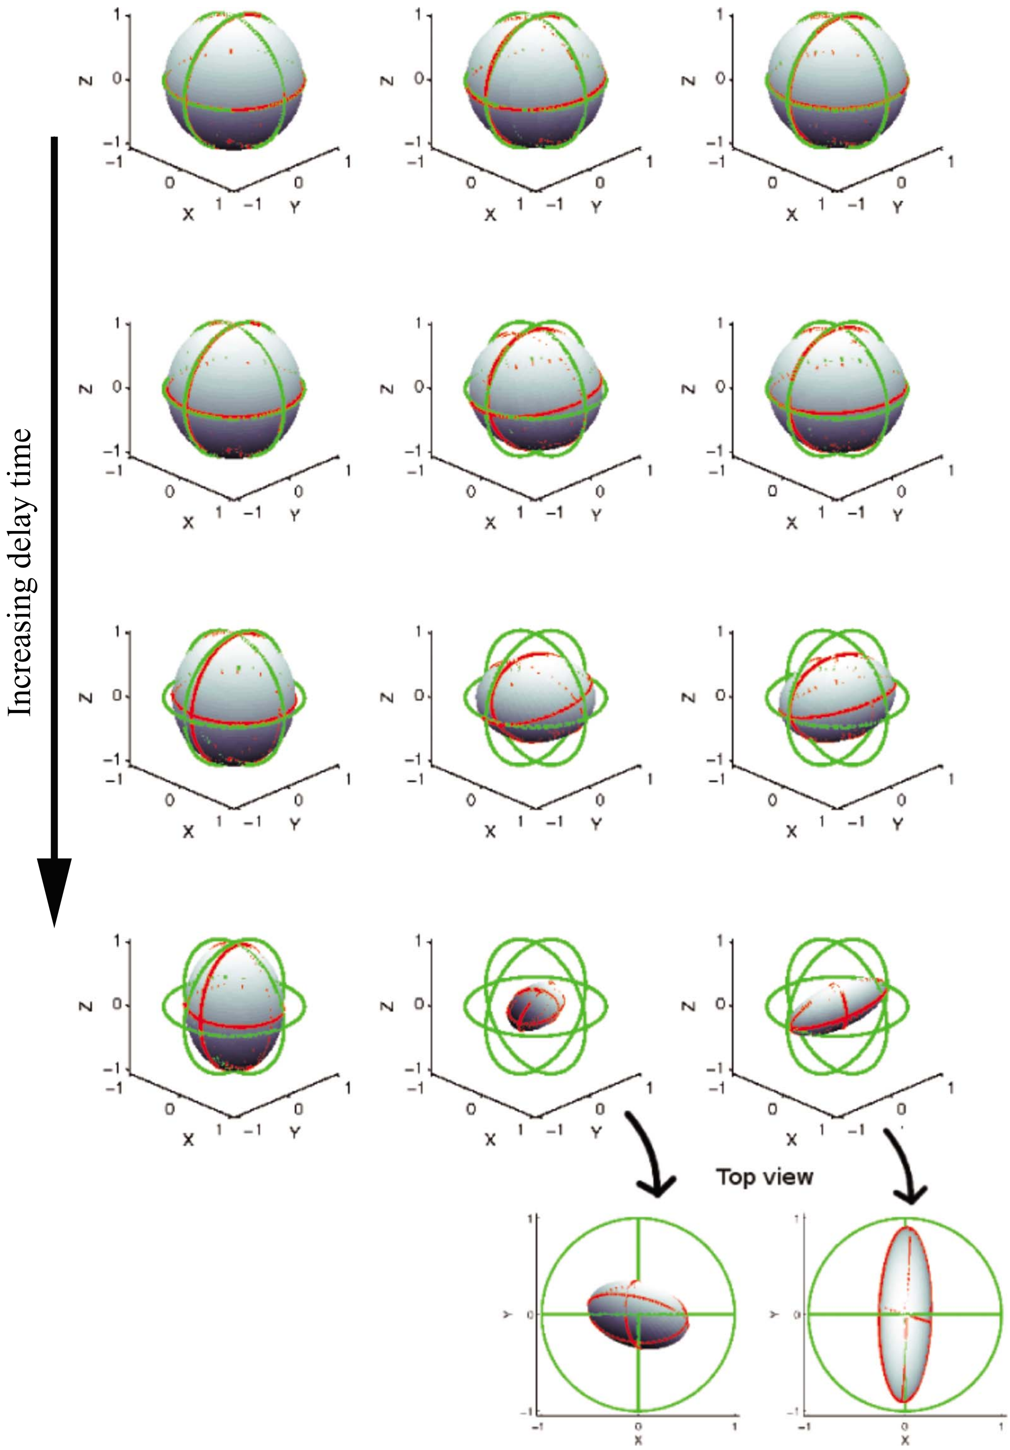

FIG. 6. (Color online) Bloch sphere transformations for three different pulse sequences on critical point ($T_2=800\pm 10$ ms). Fourier filtering was used to select a detuning range of ± 150 Hz from the Rabi frequency. Column 1: two-pulse spin echo; four different delay times moving down the column (50, 90, 170, and 330 ms). Column 2: CP_{π}, number of pulses in CP sequence equivalent to total delay time in first column ($N=11,21,41,81$). Column 3: $CP_{\pi/2}$.

is increased fractionally by 300 Hz to 8.6525 MHz. As seen in Fig. 9 (top row), the symmetric and the asymmetric CP sequences still degrade the volume fidelity compared to the simple two-pulse spin echo. However, the data show a larger volume fidelity than the model for the DDC sequences predicts (with $T_2=380$ ms), suggesting that these sequences are indeed reducing the decoherent errors, although introducing incoherence.

Increasing the sensitivity to decoherent error further by changing the Z magnetic field value to 217 G (here, $T_2=150$ ms), it is seen that the CP sequences are consistently maintaining a larger volume transformation than the two-

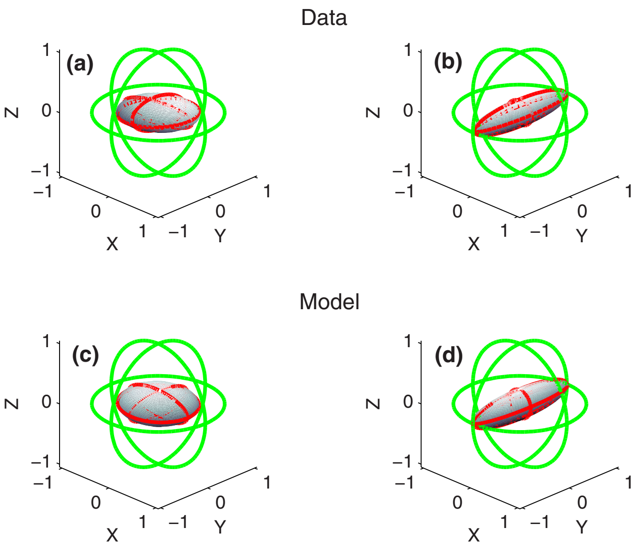

FIG. 7. (Color online) Comparison of data and model of Bloch sphere transformations for (a),(c) CP_τ and (b),(d) $CP_{\tau/2}$ pulse sequences. For this comparison, the entire frequency detuning range was used (~2.5 kHz), and for the model an rf inhomogeneity of 2% was assumed. To better compare the two sets of data, the phenomenological T_2 decay was removed from the calculation and 91 pulse iterations were performed to obtain a similar shaped transform when compared with experimental data after 61 pulses.

pulse spin echo. In this case, the incoherent processes no longer dominate, and the ions are being successfully decoupled from environmental perturbations. The model, with $T_2 = 150 \pm 10$ ms, predicts much faster volume decay for the CP sequences, and a significantly better fit is obtained by setting T_2=250 ms. This suggests that we have effectively extended the qubit lifetime by 70%.

The rotation fidelity is significantly worse in each CP scenario compared to the single-π pulse sequences; however, this is to be expected since errors in pulse area will compound when many pulses are being applied. The prevalent incoherent error source here is the inhomogeneity and stochastic fluctuations in the rf amplitude. In an attempt to reduce the sensitivity of the ions to variations in pulse area, a composite pulse sequence was used in place of the simple π pulse at zero phase. The so-called Broadband 1 (BB1) sequence consists of [24]

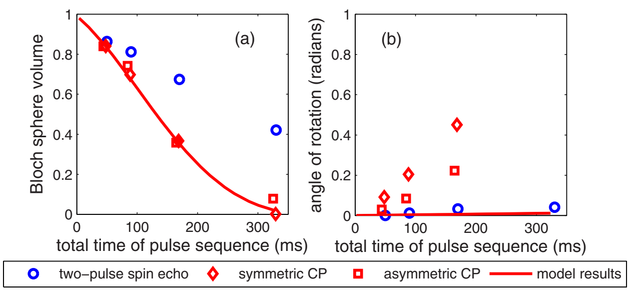

FIG. 8. (Color online) (a) Volume and (b) rotation trends for the three pulse sequences (two-pulse spin echo, CP_τ and $CP_{\tau/2}$) at the critical point. In this regime, the DDC sequences are having a detrimental effect.

$$\pi_{(phase=104.5°)} - 2\pi_{(313.4°)} - \pi_{(104.5°)} - \pi_{(0°)},$$

where the initial angle denotes the pulse area and the angle in parentheses is the phase. This was designed to be broadband with respect to the strength of the rf field and is nominally equivalent to a $\pi_{(0°)}$ pulse. Results using the BB1 sequences are also shown in Fig. 9. Using the BB1 pulse significantly improves the rotation fidelity and also increases the volume fidelity, and is thus reducing the effect of the inhomogeneity in the rf. As expected, it was seen that the effect of the BB1 pulse is practically invariant when the area is changed as a whole. However, slight errors in the relative phase and area of the constituent pulses significantly reduce its effectiveness; modeling suggests that a 1% error in both phase and intensity of one of the phase channels is enough to change BB1 from perfectly maintaining the Bloch sphere to a complete volume decay at ~300 ms. This hypersensitivity made it experimentally difficult to consistently obtain precise BB1 pulses on a day to day basis.

V. DISCUSSION

Applying the critical-point technique, where T_2 ~900 ms, the effectiveness of DDC is indeterminable since the incoherent errors, induced by technical issues such as rf inhomogeneity, completely dominates the decoherence effects. Although, operating in this regime allowed the incoherent errors to be characterized. The magnetic field was then tracked away from the critical point to increase the decoherence, while the incoherence effects remain unaffected. Altering the magnetic field and reducing T_2 to 150 ms, the DDC pulse sequences were seen to perform notably better, i.e., maintain a larger volume of the Bloch sphere transformation, than the two-pulse spin echo which removes incoherent effects. Thus, it can be concluded that the effects of decoherence are being counteracted to some extent by the dynamic decoupling sequence.

Interestingly, the symmetric and the asymmetric sequences produce contrasting Bloch sphere transformations geometrically, and for some applications choice of one over the other might offer some advantages. However, according to our fidelity definitions, the different symmetry pulse sequences are equivalent in their capacity as a quantum memory. An additional point to note is that these alternative shaped transformations are consistent with the model accounting for incoherent properties, i.e., the shape is determined exclusively by the ensemble homogeneity.

Although we place emphasis on the volume parameter as the key indication of fidelity, the rotation is also important in that it reflects the technical issues in the experiment. By reducing the technical noise sources that are inducing incoherence, we expect to increase the range of decoherence levels over which DDC can be explored. The most significant limitation currently is inhomogeneity and fluctuations in the rf field, which can be improved by optimizing the rf coil, using a thinner sample, inserting closed-loop control on the amplifier output, or designing an improved system for rf pulses and phase delays which allow for more consistent and reliable composite pulse sequences.

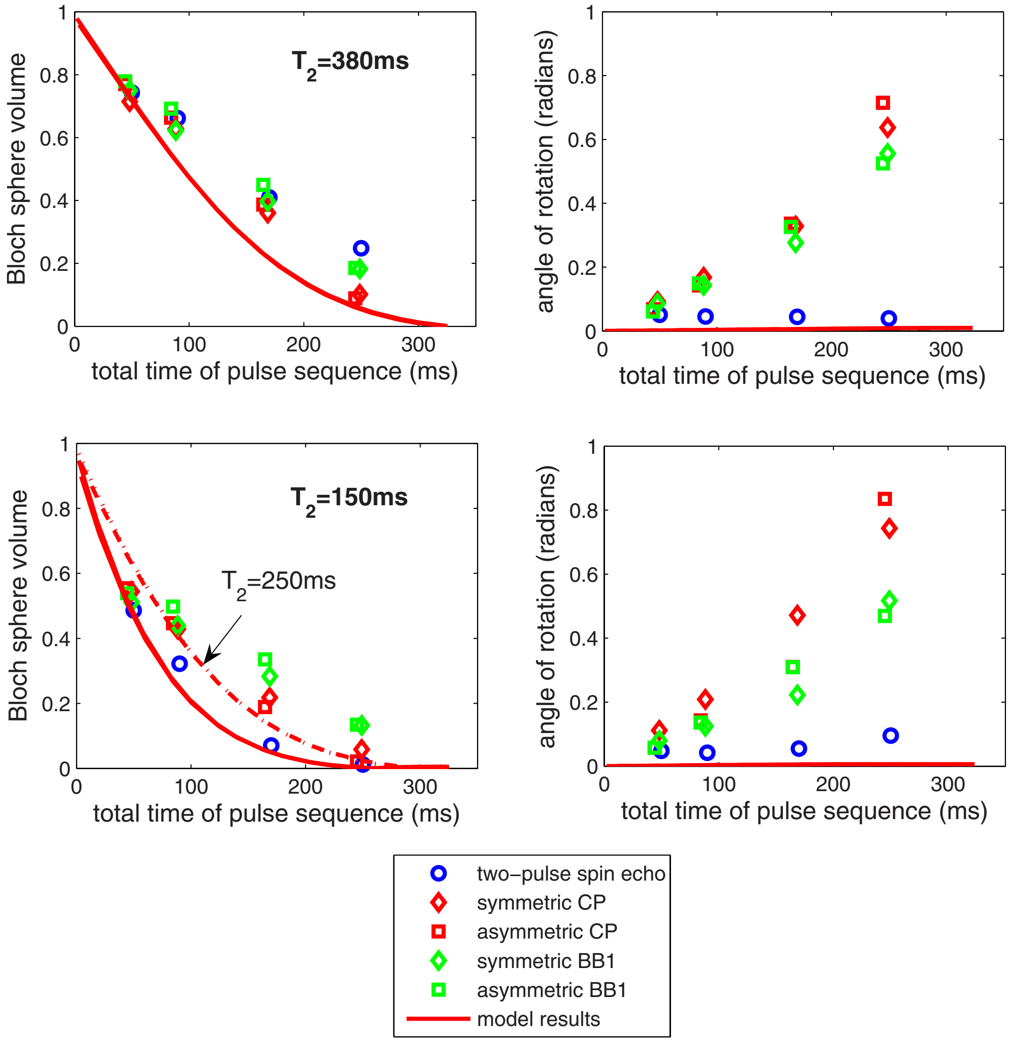

FIG. 9. (Color online) Volume and rotation fidelity results for $T_2 = 380$ ms (top row) and $T_2 = 150$ ms (bottom row). In the latter case, the Carr-Purcell sequences are maintaining a larger volume than the two-pulse spin echo, and therefore the ions are being successfully decoupled from the environmental perturbations.

The theoretical model predicts that, if the rf intensity variations of 2% can be reduced by a factor of 3, DDC will be effective at reducing the decoherence even at the critical-point field. In addition to the current application of this solid-state NMR system for investigation of DDC techniques and verification of DDC theory, being able to further extend the coherence time and the volume fidelity also adds to the appeal of this system as a quantum memory.

VI. CONCLUSION

We have experimentally demonstrated dynamic decoherence control as a means of reducing decoherent errors in a quantum ensemble. The pulse sequences used for dynamic decoupling were two forms of the Carr-Purcell pulse sequence, symmetric and asymmetric, and these were compared to a two-pulse spin echo (of equivalent total sequence duration) using the volume of the Bloch sphere transformation as a measure of fidelity. In the regime where the ensemble coherence time was measured as 150 ± 10 ms, applying DDC effectively increases this to 250 ± 40 ms.

ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the AOARD and the Australian Research Council.

[1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000).
[2] P. Zanardi and M. Rasetti, Phys. Rev. Lett. **79**, 3306 (1997).
[3] D. A. Lidar, I. L. Chuang, and K. B. Whaley, Phys. Rev. Lett. **81**, 2594 (1998).
[4] L. Viola and S. Lloyd, Phys. Rev. A **58**, 2733 (1998).
[5] E. Fraval, M. J. Sellars, and J. J. Longdell, Phys. Rev. Lett. **95**, 030506 (2005).
[6] J. J. Longdell, E. Fraval, M. J. Sellars, and N. B. Manson, Phys. Rev. Lett. **95**, 063601 (2005).
[7] A. L. Alexander, J. J. Longdell, M. J. Sellars, and N. B. Manson, Phys. Rev. Lett. **96**, 043602 (2006).
[8] J. J. Longdell, M. J. Sellars, and N. B. Manson, Phys. Rev. Lett. **93**, 130503 (2004).
[9] N. Ohlsson, R. K. Mohan, and S. Kroll, Opt. Commun. **201**, 71 (2002).
[10] M. D. Eisaman, S. Polyakov, M. Hohensee, J. Fan, P. Hemmer, and A. Migdall, Proc. SPIE **6780**, 67800K (2007).
[11] L. Viola, J. Mod. Opt. **51**, 2357 (2004).
[12] M. A. Pravia, N. Boulant, J. Emerson, A. Farid, E. M. Fortunato, T. F. Havel, and D. G. Cory, J. Chem. Phys. **119**, 9993 (2003).
[13] L. Allen and J. H. Eberly, *Optical Resonance and Two-Level Atoms* (Dover, New York, 1975).
[14] R. W. Equall, R. L. Cone, and R. M. Macfarlane, Phys. Rev. B **52**, 3963 (1995).
[15] K. Holliday, M. Croci, E. Vauthey, and U. P. Wild, Phys. Rev. B **47**, 14741 (1993).
[16] J. Longdell, Ph.D. thesis, Australian National University, 2003.
[17] G. Liu and B. Jacquier, *Spectroscopic Properties of Rare Earths in Optical Materials* (Springer, New York, 2005).
[18] E. Fraval, M. J. Sellars, and J. J. Longdell, Phys. Rev. Lett. **92**, 077601 (2004).
[19] J. Mlynek, N. C. Wong, R. G. DeVoe, E. S. Kintzer, and R. G. Brewer, Phys. Rev. Lett. **50**, 993 (1983).
[20] Additional shots were also taken for normalization and to monitor changes in T_2 over the course of the measurements. The duration of a single "Bloch sphere" measurement was typically ~ 1 s.
[21] Linearity of the system is a consequence of negligible coupling between Pr ions.
[22] The model does predict some overall rotation arising from a combination of detunings and rf inhomogeneity; however, this is negligible when compared to the rotation observed in the experimental data, as will be seen in Figs. 8 and 9.
[23] R. Macfarlane, J. Lumin. **45**, 1 (1990).
[24] S. Wimperis, J. Magn. Reson., Ser. A **109**, 221 (1994).

APPENDIX B

Spontaneous emission from an ensemble of two-level atoms

In this Appendix the solution is derived for spontaneous emission from an ensemble of two-level atoms, following an approach similar to Weisskopf and Wigner,[1] who solved the theoretical problem of spontaneous emission for the first time in 1930 [131], applying the quantum electrodynamic theory developed by Dirac. This solution is applied in Chapter 4 to determine the state of the atomic system after a single photon has been detected from an ensemble of atoms initially in their excited state.

B.1 The problem

The system consists of N two-level atoms, interacting with all modes of the electromagnetic field, where it is assumed in quantizing the field that the system is enclosed in a box of volume V. The Hamiltonian in the interaction picture, after applying the rotating-wave approximation, is found in §4.2 to be;

$$\mathscr{V} = \sum_j \sum_k \hbar g_k \left(\hat{\sigma}_j^\dagger \hat{a}_k e^{i\mathbf{k}.\mathbf{r}_j} e^{i\Delta_{jk}t} + \hat{\sigma}_j \hat{a}_k^\dagger e^{-i\mathbf{k}.\mathbf{r}_j} e^{-i\Delta_{jk}t} \right) \tag{B.1}$$

Where j indexes the atom number, k the field mode, and the remaining parameters are;

[1]The assumption that the approach presented here is consistent with the original work of Weisskopf and Wigner is indirectly implied by following the texts [36], [130] and [107], since the author could not read German well enough to verify directly.

$$
\begin{aligned}
g_k &= \text{coupling between a single atom and the } k^{\text{th}} \text{ field mode} \\
\hat{\sigma}_j^\dagger &= \text{raising operator for } j^{\text{th}} \text{ atom} \\
\hat{\sigma}_j &= \text{lowering operator for } j^{\text{th}} \text{ atom} \\
\hat{a}_k^\dagger &= \text{raising operator for } k^{\text{th}} \text{ field mode} \\
\hat{a}_k &= \text{lowering operator for } k^{\text{th}} \text{ field mode} \\
\mathbf{k} &= \text{wavevector of the } k^{\text{th}} \text{ field mode} \\
\mathbf{r}_j &= \text{position vector of the } j^{\text{th}} \text{ atom} \\
\Delta_{jk} &= \omega_j - \nu_k \text{ , the detuning of the } j^{\text{th}} \text{ atom from the } k^{\text{th}} \text{ field mode} \\
t &= \text{time}
\end{aligned}
$$

To determine the dynamics of the system described by this interaction Hamiltonian, we will solve the Schrödinger equation for the probability amplitudes of an arbitrary state vector. This is an extension of the approach presented in §6.3 of reference [36] for a single atom. An arbitrary state of the atom-field system can be written;

$$
|\psi(t)\rangle = \alpha(t)\,|e_N, 0\rangle + \sum_j \sum_k \beta_{kj}(t)\,|g_j, 1_k\rangle + \ldots \tag{4.9}
$$

where $|e_N, 0\rangle \rightarrow$ All N atoms are in excited state $|e\rangle$ and there are 0 photons

$|g_j, 1_k\rangle \rightarrow$ The j^{th} atom is in ground state $|g\rangle$, and there is 1 photon in mode k

In order to solve the Schrödinger equation, it is assumed that the time t is short compared to the excited state lifetime. In this case only the first two terms in the equation 4.9 are retained, where there is either none, or one excitation of the field.

B.2 Schrödinger equation

Substitution of the state vector (equation 4.9) into the Schrödinger equation (equation 3.19), yields a set of coupled differential equations for the probability amplitudes $\alpha(t)$ and $\beta_{kj}(t)$;

$$
\frac{\mathrm{d}}{\mathrm{d}t}\alpha(t) = -i\sum_j \sum_k g_k e^{i\mathbf{k}.\mathbf{r}_j} e^{i\Delta_{jk}t} \beta_{kj}(t) \tag{B.2}
$$

$$
\frac{\mathrm{d}}{\mathrm{d}t}\beta_{kj}(t) = -i g_k^* e^{-i\mathbf{k}.\mathbf{r}_j} e^{-i\Delta_{jk}t} \alpha(t) \tag{B.3}
$$

To obtain an equation dependent on $\alpha(t)$ alone, equation B.3 is integrated, with the initial condition given by $|\psi(0)\rangle = |e_N, 0\rangle$;

$$\beta_{kj}(t) = -ig_k^* e^{-i\mathbf{k}.\mathbf{r}_j} \int_0^t \alpha(\tau) e^{-i\Delta_{jk}\tau} \, \mathrm{d}\tau \tag{B.4}$$

Then substitution for $\beta_{kj}(t)$ into equation B.2 gives the following integro-differential equation for $\alpha(t)$;

$$\frac{\mathrm{d}}{\mathrm{d}t}\alpha(t) = -\sum_j \sum_k |g_k|^2 \int_0^t \alpha(\tau) e^{i\Delta_{jk}(t-\tau)} \, \mathrm{d}\tau \tag{B.5}$$

Provided the field-quantization volume $V = L^3$ is large, then the field states are closely spaced in frequency and the summation over k can be replaced by an integral.

$$\sum_k |g_k|^2 = \frac{V}{(2\pi)^3} \int_0^{2\pi} \mathrm{d}\phi \int_0^{\pi} \sin\theta \, \mathrm{d}\theta \int_0^{\infty} k^2 |g_k|^2 \, \mathrm{d}k \tag{B.6}$$

In the spherical-coordinate frame, if the field wavevector $\mathbf{k}$ is along the $\hat{z}$-axis, and the atomic dipole vector $\hat{\mathbf{d}}$ is along the $\hat{y}$-axis, then from equation 4.6;

$$g_k = \sqrt{\frac{\nu_k}{2\epsilon_0 \hbar V}} \, \mu \, \hat{\mathbf{d}}.\hat{\boldsymbol{\epsilon}} = \sqrt{\frac{\nu_k}{2\epsilon_0 \hbar V}} \, \mu \, \sin\theta$$

Where θ is the angle between the $\hat{x}$-axis and the field polarization $\hat{\boldsymbol{\epsilon}}$. Also substituting $k = \nu_k / c$ into equation B.6 gives;

$$\begin{aligned}\sum_k |g_k|^2 &= \frac{V}{(2\pi)^3} \frac{\mu^2}{2\hbar\epsilon_0 V} \int_0^{2\pi} \mathrm{d}\phi \int_0^{\pi} \sin^3\theta \, \mathrm{d}\theta \int_0^{\infty} \frac{\nu_k^3}{c^3} \, \mathrm{d}\nu_k \\ &= \frac{\mu^2}{6\pi^2 \hbar \epsilon_0 c^3} \int_0^{\infty} \nu_k^3 \, \mathrm{d}\nu_k\end{aligned}$$

Then equation B.5 becomes;

$$\frac{\mathrm{d}}{\mathrm{d}t}\alpha(t) = -\frac{\mu^2}{6\pi^2 \hbar \epsilon_0 c^3} \sum_j \int_0^{\infty} \nu_k^3 \, \mathrm{d}\nu_k \int_0^t \alpha(\tau) e^{i\Delta_{jk}(t-\tau)} \, \mathrm{d}\tau \tag{B.7}$$

B.3 Solution for probability amplitudes $\alpha(t)$ and $\beta_{kj}(t)$

Equation B.7 is a formidable equation which cannot be solved for the general case. However an approximate solution can be obtained by making a simplifying assumption. Due to the detunings Δ_{jk}, The integral of the $e^{i\Delta_{jk}(t-\tau)}$ term over frequencies ν_k will only be appreciably different from zero when $t \approx \tau$, i.e. before the inhomogeneity in frequency causes dephasing of the oscillations. Therefore, the next approximation to be made is to let $\alpha(\tau) = \alpha(t)$, which can then be taken outside of the integral. Another way of interpreting this approximation, is in saying that the system has no "memory" of earlier times. The integral over τ is then determined as;

$$\int_0^t e^{i\Delta_{jk}(t-\tau)}\ \mathrm{d}\tau = \pi\delta\left(\Delta_{jk}\right) + i\mathcal{P}\left(\frac{1}{\Delta_{jk}}\right)$$

Where δ is the Dirac delta function, and $\mathcal{P}$ denotes the 'principle value' integral [132]. This is a purely imaginary term, and physically indicates the dressing of the atomic states by the continuum of vacuum modes. From here on it is assumed that the atomic energy levels are re-defined to absorb this Stark-shift, and thus the $i\mathcal{P}\left(\frac{1}{\Delta_{jk}}\right)$ term can be omitted. Finally, equation B.7 becomes;

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t}\alpha(t) &= -\alpha(t)\ \frac{\mu^2}{6\pi^2\hbar\epsilon_0 c^3}\sum_j \pi \int_0^\infty \nu_k^3\ \delta\left(\omega_j - \nu_k\right)\ \mathrm{d}\nu_k \\
&= -\alpha(t)\ \frac{\mu^2}{6\pi\hbar\epsilon_0 c^3}\sum_j \omega_j^3 \\
&= -\gamma\ \alpha(t) \qquad \text{where} \quad \gamma = \frac{\mu^2}{6\pi\hbar\epsilon_0 c^3}\sum_j \omega_j^3
\end{aligned} \tag{B.8}$$

This is easily solved to give;

$$\alpha(t) = e^{-\gamma t} \tag{B.9}$$

The solution for $\beta_{kj}(t)$ is obtained by substituting $\alpha(t) = e^{-\gamma t}$ into equation B.4;

$$\begin{aligned}
\beta_{kj}(t) &= -ig_k e^{-i\mathbf{k}.\mathbf{r}_j}\int_0^t e^{-\left(\gamma + i\Delta_{jk}\right)\tau}\ \mathrm{d}\tau \\
&= g_k e^{-i\mathbf{k}.\mathbf{r}_j}\frac{1 - e^{-\left(\gamma + i\Delta_{jk}\right)t}}{i\gamma - \Delta_{jk}}
\end{aligned} \tag{B.10}$$

We finally arrive at the full solution for the state at time t by combining equations 4.9, B.9 and B.10;

$$|\psi(t)\rangle = e^{-\gamma t}\left|e_N, 0\right\rangle + \sum_j\sum_k g_k e^{-i\mathbf{k}.\mathbf{r}_j}\frac{1 - e^{-\left(\gamma + i\Delta_{jk}\right)t}}{i\gamma - \Delta_{jk}}\left|g_j, 1_k\right\rangle \tag{B.11}$$

B.4 Properties of the decay rate γ

In a completely isolated system, quantum mechanics tells us that an eigenvalue takes an exact value. In an atomic system, the presence of infinitely many field modes to which the atom is coupled has the effect of causing decay of excited states, and as a result the eigenvalue of the transition frequency has a finite width quantified by γ. The decay time T_1, introduced in Chapter 2 as a measured value for a given system, is simply the inverse of γ. A long T_1 can be obtained in situations where the atomic system is only weakly coupled to the modes of the electromagnetic field.

To understand the behaviour of the decay rate γ, we consider equation B.8 for a few simple cases. Firstly, for a single atom, $\sum_j \omega_j^3 \to \omega^3$, and the value of γ reduces to;

$$\gamma = \frac{\omega^3 \mu^2}{6\pi\hbar\epsilon_0 c^3} \equiv \frac{\Gamma}{2} \tag{B.12}$$

where Γ is the Weisskopf-Wigner spontaneous emission rate. Note that Γ is also precisely the rate of population decay that is obtained using Fermi's Golden Rule, derived using first-order perturbation theory, and therefore valid only for short times [130]. Fermi's Golden Rule is found from the Wiesskopf-Wigner solution by taking the small time limit $e^{-\frac{\Gamma}{2}t} \approx 1 + \frac{-\Gamma}{2}t$. Noteworthy features of Γ are that it is dependent on the transition frequency cubed, and the transition dipole moment squared.

As a second example, consider an ensemble of N atoms, with no inhomogeneity in the atomic transition, then $\sum_j \omega_j^3 \to N\omega_0^3$ and the differential equation becomes;

$$\alpha(t) = e^{-\gamma t} \qquad \text{where} \quad \gamma = N\frac{\omega_0^3 \mu^2}{6\pi\hbar\epsilon_0 c^3} \tag{B.13}$$

This is a very interesting result; the rate of decay from the excited state is N times faster for an ensemble of N identical atoms as compared to a single atom. One way of interpreting this is to incorporate the additional factor of N into the definition of the coupling coefficient g_k. Then this is simply saying that the coupling of N atoms to a mode of the field is increased by a factor of N compared to the coupling of a single atom.

In considering the typical number of atoms involved in ensemble experiments, in excess of 10^{10} for the experiments presented in this thesis, it would appear from equation B.13 that the dynamical behaviour of the ensemble will be significantly different to that of a single atom. However this is not necessarily true, and the reason for this lies with the assumption of the coupling coefficient g_k being independent of any spatial properties of the atomic ensemble. As an example, consider a cylindrical ensemble which is much longer than it is wide. In this case the field modes which lie along the axis of the cylinder will be enhanced by a factor close to N, whereas a field mode perpendicular to the ensemble axis has only a small fraction of the atoms contributing in a comparable solid angle. This enhanced coupling is closely linked to the optical depth along a given direction $\mathbf{k}$.

The increased decay rate for an ensemble of atoms is one result of an effect referred to as *superradiance*. In his seminal paper of 1954, Dicke considered the nature of the spontaneous emission from many atoms, and showed that not only is the transition broadened, as we have shown here, but that the intensity of the emission has a component proportional to N^2 [110]. This may seem surprising; that the inherently incoherent process of spontaneous emission can somehow develop into *coherent* emission, however it is essentially stimulated emission which is triggered by a spontaneous event, and if the gain is large then the build up of coherence makes perfect sense.[2]

For a more general discussion of superradiance, including not only the temporal properties, but the spatial distribution of radiated intensity, the influence of the ensemble geometry, and the effect of inhomogeneous broadening, see Chapter 8 of Allen and Eberley [94], and references [97], [133]. Some properties of superradiance in a rare-earth ion system were reported in [134].

[2]With the addition of feedback, this is the operating principle of a laser.

APPENDIX C

Heterodyne phase measurement

This section describes the experimental methods and numerical algorithms used to determine the phase of heterodyne signals at various frequencies. These methods were applied in characterizing the phase properties of the four-level echo sequence in Section 6.5.5.

In the experimental setup, the local oscillator and the signal beams travel along independent paths with many optical components involved, before being superimposed on a beam-splitter for the heterodyne detection. Mechanical vibrations throughout causes the phase relationship between these two beams to vary randomly from shot to shot. Extracting phase information from the heterodyne signal therefore involves first measuring the phase of the interferometer (phase difference between local oscillator and signal beams) in each shot.

An example scope screen-shot of a pulse sequence for phase-sensitive detection is shown in figure C.1. The basic idea is to have optical pulses along the signal beam path which can be used to determine the interferometer phase. These reference pulses are transmitted through a spectral trench in the sample's absorption profile, as distant in frequency as possible from any absorbing ions so as to minimize dispersion.

The oscillators clocking the various RF sources used to drive the AOMs were synchronized, meaning that the relative difference in optical frequencies between different beam paths in the absence of noise was known precisely (see Appendix F for instrument specifications). The digital oscilliscope used to record the heterodyne signal however was unable to be synchronized to the experiment RF sources. As a result, to accurately measure the phase of a given signal, it is necessary to also measure the frequency difference between the data-acquisition and the experiment oscillators. The optimal means for achieving this is to have a reference pulse at the beginning of the time-frame and the same again at the end of the frame.

The algorithm for determining the frequency offset between the acquisition and experiment clocks begins by numerically heterodyning the signal at the expected beat frequency of one set of the reference pulses, i.e. multiply the signal by $e^{i\omega_{\mathrm{ref}}t}$. Then the phase of each of these reference pulses, $\phi_{\mathrm{ref}}^{t_1}$ and $\phi_{\mathrm{ref}}^{t_2}$ for the first and second pulse respectively, is determined using the ratio of real and imaginary components of the time-domain signal centred on times t_1 and t_2 (fig. C.1). This phase measurement is made for all shots, and

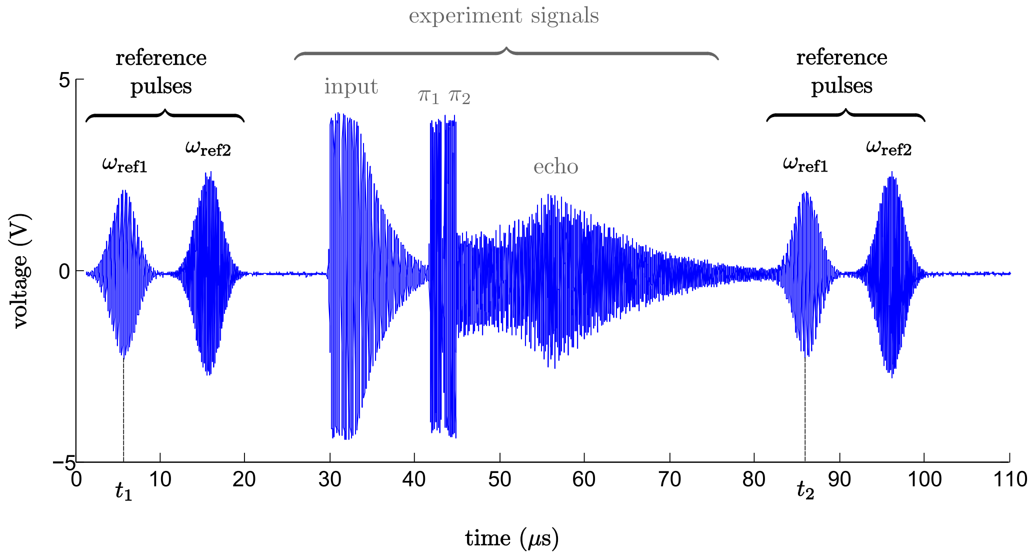

Figure C.1: An example recording of a single-shot of a four-level echo pulse sequence with co-propagating beam geometry using the balanced heterodyne detector. The pulses for the 4LE are contained in the domain labelled 'experiment signals', and before and after this region there are pairs of reference pulses. Recall that the frequencies of the input, π_1, π_2 and echo signals are all different (see Chapter 6), and the frequencies ω_{ref1} and ω_{ref2} are different again. Having reference pulses at both ends of the shot allows for maximum sensitivity in measuring the frequency offset between the detection and experiment systems, and two frequencies are required so as to precisely determine the time when the phase of all the different frequencies (determined by the programming of the RF source) is equal.

the frequency offset is calculated using the average phase change;

$$\Delta\omega_{\mathrm{ref}} = \frac{\left\langle \phi_{\mathrm{ref}}^{t_1} - \phi_{\mathrm{ref}}^{t_2} \right\rangle_{\text{all shots}}}{t_2 - t_1}$$

The ratio of this frequency offset to the absolute frequency, $\Delta\omega_{\mathrm{ref}}/\omega_{\mathrm{ref}}$ gives the fractional offset between the acquisition and experiment clocks, and can be used to determine the offset for an arbitrary frequency pulse. Typically, $\Delta\omega_{\mathrm{ref}}/\omega_{\mathrm{ref}}$ was on the order of 10^{-4}, or 1.2 kHz in a 10 MHz detected beat signal, and this would drift significantly over the course of hours.

As mentioned above, mechanical vibration of optical components along the different beam paths adds random Doppler-shifts to the light frequencies. In addition to the randomized phase of the local oscillator relative to the signal beam between shots, which are typically 200 ms apart, there is also a measurable phase gradient across the $\sim 100\ \mu$s window of a single shot. This phase change, given by $\left(\phi_{\mathrm{ref}}^{t_1} - \phi_{\mathrm{ref}}^{t_2}\right) - \left\langle \phi_{\mathrm{ref}}^{t_1} - \phi_{\mathrm{ref}}^{t_2} \right\rangle_{\text{all shots}}$, was found to be at worst $\pm 2°$, and so was corrected for to first order by assuming a linear variation in phase across each shot.

With all the above procedures, it would be possible to measure the phase of a signal at frequency ω_{ref}, however to determine phase of different frequency signals, the phase relationship between the different frequencies as generated by the RF source is required, i.e. the point in time when the phase of all frequencies were equal. For this purpose, a second frequency reference pulse is used, and the 'phase-reset' time is found as the time when the phases of these two frequency signals are equal.[1]

For the 4LE measurements described in §6.5.5, the heterodyne beat frequencies of the reference pulses were 4.2 MHz and 19.0 MHz, spectrally positioned in the side of the trenches for the input pulse and echo. In the time domain, the pulses were 5 μs fwhm in length, and the separation $(t_2 - t_1)$ was 80 μs. The accuracy of phase measurements with this configuration of reference pulses was typically 0.044 radians, or 2.5°, at the echo beat frequency of 18.6 MHz.

[1]If it were accurate enough, the oscilloscope trigger could be used to mark the phase-reset time, however the trigger was seen to exhibit jitter of approximately 3ns, which is the order of 2π phase for beat signal in the 10MHz range.

APPENDIX D

Photon Statistics

Examining statistics of a light field can yield invaluable information about the source of radiation. Also, it is in the statistical analysis where the subtly differing predictions of quantum and classical descriptions of the field can be tested.

In the initial section of this appendix, §D.1, the concept of correlation functions as a means of describing the statistics of a fluctuating field is introduced. Following this, the statistical properties for three different types of fields; coherent-, incoherent- and number-states are considered in sections D.2, D.3 and D.4 respectively.

This appendix provides a theoretical background for Chapter 5, where the statistics of the photon pairs generated in RASE are considered. Also, the analysis of experimental results of Chapter 7 includes aspects of the theory presented in §D.3 on the dynamics of photon bunching.

D.1 Correlation functions and coherence

As discussed in the pioneering work of Glauber, the physical process of detecting an electromagnetic field usually relies on absorption of photons [114]. As such, the probability for a detector absorbing a photon at position $\mathbf{r}$ and time t is related to $|\langle f|\,\hat{a}(\mathbf{r},t)\,|i\rangle|^2$, where $\hat{a}(\mathbf{r},t)$ is the field annihilation operator in the Heisenberg picture, and $|i\rangle$ and $|f\rangle$ are the initial and final states of the field respectively. The total counting rate, or intensity, that is recorded by the detector is related to the sum of transition probabilities over all possible final states:

$$\mathcal{I}(\mathbf{r},t) = \mathcal{I}_0 \sum_f \langle i|\,\hat{a}^\dagger(\mathbf{r},t)\,|f\rangle\,\langle f|\,\hat{a}(\mathbf{r},t)\,|i\rangle = \mathcal{I}_0\,\langle i|\,\hat{a}^\dagger(\mathbf{r},t)\hat{a}(\mathbf{r},t)\,|i\rangle \equiv G^{(1)}$$

where $G^{(1)}$ is defined as the first-order correlation function, $\hat{a}^\dagger(\mathbf{r},t)$ is the field creation operator, $\mathcal{I}_0$ is a constant relating photon number to intensity, and the closure relation $\sum_f |f\rangle\,\langle f| = 1$ has been used.

The measurements made in the famous work of Hanbury-Brown and Twiss in measuring the angular extent of the star Sirius involved coincidences of detection events between two different space-time points ($\mathbf{r}_1$,t_1) and ($\mathbf{r}_2$,t_2) [135, 136]. The measurement made in this 'intensity interferometer' can be mathematically constructed similarly to above, where this time the initial state is subject to two annihilation events at two different points in space-time. The resulting coincidence detection is defined by the second-order correlation function;

$$G^{(2)} = \mathcal{I}_0^2 \langle i| \hat{a}^\dagger(\mathbf{r}_2,t_2)\hat{a}^\dagger(\mathbf{r}_1,t_1)\hat{a}(\mathbf{r}_1,t_1)\hat{a}(\mathbf{r}_2,t_2)|i\rangle$$

Although the practicality of such a measurement is questionable, one could in theory measure detection coincidences between an arbitrary number of space-time points. In general, the n^{th} order correlation function is;

$$G^{(n)} \propto \langle i| \hat{a}^\dagger(\mathbf{r}_n,t_n)\ldots\hat{a}^\dagger(\mathbf{r}_2,t_2)\hat{a}^\dagger(\mathbf{r}_1,t_1) \times \hat{a}(\mathbf{r}_1,t_1)\hat{a}(\mathbf{r}_2,t_2)\ldots\hat{a}(\mathbf{r}_n,t_n)|i\rangle$$

This family of correlation functions defines a means of quantifying the properties of fluctuating electromagnetic fields. As an example, the simplest manifestations of correlations in optical fields are well-known interference effects. Spatial interference fringes (as seen in Young's double-slit experiment) or the temporal coherence of a laser beam (parameterized by the 'coherence time') are connected to the properties of the first-order correlation function $G^{(1)}$.[1] This first-order correlation function however does not carry any signature of field quantization; all possible behaviours of $G^{(1)}$ are explicable by modelling the field as a wave.

In contrast, the second-order correlation function does behave differently depending on the model of the field. For this reason, the second-order correlation function has become an important tool in quantum-optics experiments. In normalized form, this function is re-written as;

$$g^{(2)} = \frac{\left\langle \hat{a}_2^\dagger\hat{a}_1^\dagger\hat{a}_1\hat{a}_2 \right\rangle}{\left\langle \hat{a}_1^\dagger\hat{a}_1 \right\rangle \left\langle \hat{a}_2^\dagger\hat{a}_2 \right\rangle} = \frac{\langle : \hat{n}_1\hat{n}_2 : \rangle}{\langle \hat{n}_1\rangle \langle \hat{n}_2\rangle} \tag{D.1}$$

This has been generalized to an ensemble-average, denoted by $\langle\rangle$. The subscripts 1 and 2 denote different space-time points. The number-operator $\hat{n}$ is equal to $\hat{a}^\dagger\hat{a}$ and the :: indicates 'normal-ordering' of the operators, where creation operators appear to the left of annihilation operators in the product.

The essential difference between quantum and classical theories is that in a quantum model, two field operators do not, in general, commute with one another (hence the need to specify normal-ordering of operators in equation D.1). To see this explicitly, consider the second-order *auto-correlation* function for a given field, i.e. photon-detection coincidences measured at the *same* point in space and time (for example measured by a number-resolving detector);

[1] See Chapter 4 of Mandel and Wolf ([28]) for an introductory discussion of properties described by field correlations.

$$g^{(2)}(0) = \frac{\langle \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} \rangle}{\langle \hat{a}^\dagger \hat{a} \rangle \langle \hat{a}^\dagger \hat{a} \rangle} = \frac{\langle \hat{n}^2 \rangle - \langle \hat{n} \rangle}{\langle \hat{n} \rangle^2}$$
$$= \frac{\langle \hat{n} \rangle^2 + \mathrm{Var}\,(n) - \langle \hat{n} \rangle}{\langle \hat{n} \rangle^2} \tag{D.2}$$

where the boson commutation relation $[\hat{a}, \hat{a}^\dagger] = 1$ has been used, and $\mathrm{Var}\,(n) = \langle n^2 \rangle - \langle n \rangle^2$ is the variance of the photon-number distribution. Alternatively, for two independent fields (subscripts x and y) with commuting operators, $g^{(2)}$ is;

$$g^{(2)}_{xy}(0) = \frac{\left\langle \hat{a}^\dagger_y \hat{a}^\dagger_x \hat{a}_x \hat{a}_y \right\rangle}{\left\langle \hat{a}^\dagger_y \hat{a}_y \right\rangle \left\langle \hat{a}^\dagger_x \hat{a}_x \right\rangle} = \frac{\langle \hat{n}_x \hat{n}_y \rangle}{\langle \hat{n}_x \rangle \langle \hat{n}_y \rangle} \tag{D.3}$$

In a *classical* model, where all operators commute, then the auto-correlation function of equation D.2 would be similar in form to equation D.3, i.e. $g^{(2)}(0) = \langle \hat{n}^2 \rangle / \langle \hat{n} \rangle^2$. Therefore the second-order auto-correlation function can show different results depending on the model used to describe the field. A field that has a measured value of $g^{(2)}$ that is *not* possible assuming a classical field model is said to be 'non-classical' in nature.

In the following sections, the photon statistics obtained from different sources of radiation are considered, along with the properties of the second-order auto-correlation function $g^{(2)}$.

D.2 Coherent source

In a classical picture, light is considered as an electromagnetic wave, and the most stable type is described by a wave with constant frequency (ω), phase (ϕ) and amplitude ($\mathcal{E}_0$), written as a function of distance x and time t as:

$$\mathcal{E}(x,t) = \mathcal{E}_0 \sin\left(kx - \omega t + \phi\right) \tag{D.4}$$

The light emitted by an ideal single-mode laser is a good approximation of this field.

Classically we would expect such a field to have a perfectly constant intensity. However, if the *detection* process is considered quantized (a semiclassical model), then there are random fluctuations in the recorded intensity. For example in a photomultiplier tube, the precise time at which the field liberates individual electrons via the photoelectric effect will fluctuate randomly.

The state analogous to this perfectly stable wave in the fully quantum-mechanical picture is referred to as a coherent state. In this model where the field is quantized, fluctuations can be thought of as a result of the random arrival times of individual photons at the detector. However the *mean* photon flux is constant.

The probability of detecting n photons in a state described with a mean photon number of $\langle n \rangle$ is given by;[2]

$$\mathcal{P}(n) = \frac{\langle n \rangle^n \, e^{-\langle n \rangle}}{n!}, \qquad n = 0, 1, 2, \ldots \tag{D.5}$$

The noise inherent to this probability distribution can be characterized by the variance;

$$\text{Var}\,(n) = \sum_{n=0}^{\infty} (n - \langle n \rangle)^2 \, \mathcal{P}(n) = \langle n \rangle \tag{D.6}$$

In statistics this is known as a poissonian distribution.[3] In quantum physics, these states were first discovered by Schrödinger as the solutions of a quantum-oscillator having minimum uncertainty product [137, 138]. Glauber formed a completely quantum-theoretic definition of coherence, and coined the term 'coherent state' as being a state in which the normalized correlation-function for all orders $g^{(n)}$ is unity [111, 114].

Mathematically, the coherent state $|\alpha\rangle$ is expanded in terms of number states $|n\rangle$ as;

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} \, |n\rangle \tag{D.7}$$

This state is an eigenstate of the annihilation operator $\hat{a}$, since $\hat{a}\,|\alpha\rangle = \alpha\,|\alpha\rangle$, and the mean photon number is $\langle \hat{n} \rangle = |\alpha|^2$. In the limit of large mean photon-number, this state approaches the classical field of equation D.4.

Since the statistics describing the most stable type of classical light that can be envisaged (a perfectly coherent beam of constant intensity) are described by a poissonian distribution, then this sets a benchmark for the minimum possible noise classically achievable.

The second order auto-correlation function for a coherent field is found by substitution of $\text{Var}\,(n) = \langle n \rangle$ into equation D.2;

$$g^{(2)} = 1 \tag{D.8}$$

D.3 Incoherent source

The most familiar sources of light that surround us are typically incoherent; one wave-packet of the field is unrelated in phase to any other. Examples include the radiation emitted by a thermal body like a star, the humble incandescent bulb, or spontaneous emission from an ensemble of atoms.

To determine the statistical properties of incoherent light, we model the field as a superposition of many coherent states [111, 139]. The phase of each of these individual coherent states is randomized. The density matrix of this state can be written as;

[2] For a general introduction, see [27, 121], and for a more detailed mathematical approach see [28].

[3] In general, Poissonian statistics applies to independent random processes which have a discrete set of possible outcomes

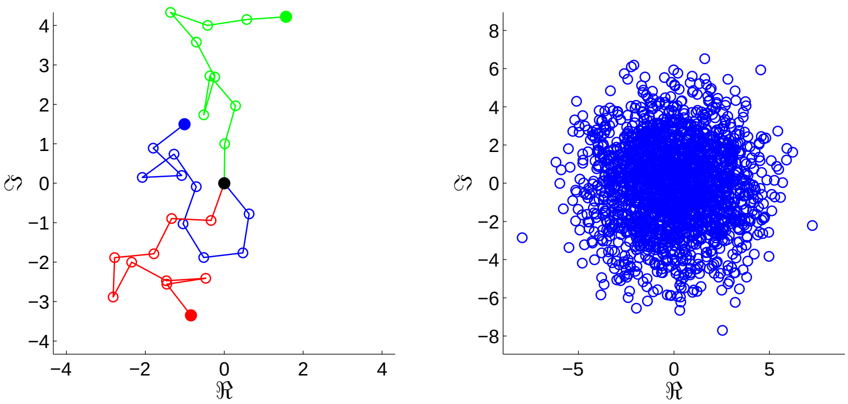

Figure D.1: The total amplitude of an incoherent state can be considered as a complex number that undergoes a random walk. The graph on the left depicts three random walks in the complex plane. Each path consists of ten steps (i.e. $j = 10$) of unit length and randomized phase. On the right hand side, the final positions of 2000 random walks are shown. In the limit of a large number of steps, this approaches a gaussian distribution.

$$\rho = |\alpha_1 + \alpha_2 + \ldots + \alpha_j\rangle \langle \alpha_1 + \alpha_2 + \ldots + \alpha_j| \tag{D.9}$$

We can think of the amplitudes α_n as being complex numbers with random phases. The sum over all α_n can then be interpreted as a random walk in the complex plane. An example of such a random walk for $j = 10$ is shown in figure D.1, along with a plot of the resultant location of this random walk if it is repeated many times. In the limit $j \to \infty$, the probability that the resultant of this sum is α approaches a gaussian distribution of the form;

$$\mathcal{P}(\alpha) = \frac{1}{\pi \langle n \rangle} e^{-|\alpha|^2/\langle n \rangle} \tag{D.10}$$

where $\langle n \rangle = \sum_n |\alpha_n|^2$ is the mean number of quanta in the mode.

Using this probability distribution, the state ρ can be re-written as a statistical mixture in the coherent state basis (referred to as the Glauber-Sudarshan $\mathcal{P}$-representation [111, 140]);

$$\rho = \int \mathcal{P}(\alpha) |\alpha\rangle \langle \alpha| \, \mathrm{d}^2\alpha = \frac{1}{\pi \langle \mathrm{n} \rangle} \int \mathrm{e}^{-|\alpha|^2/\langle \mathrm{n} \rangle} |\alpha\rangle \langle \alpha| \, \mathrm{d}^2\alpha \tag{D.11}$$

where $\mathrm{d}^2\alpha$ is equal to d[real(α)] $\times$ d[imag(α)]. This is in turn converted into the number-state basis using equation D.7;

$$\rho = \sum_m \frac{\langle n \rangle^m}{(1 + \langle n \rangle)^{1+m}} |m\rangle \langle m| \tag{D.12}$$

Thus the probability distribution in the number-state basis for this incoherent field is;

$$\mathcal{P}(n) = \frac{\langle n \rangle^m}{(1 + \langle n \rangle)^{1+m}} \tag{D.13}$$

This is known in probability theory as the *geometric distribution* since $\mathcal{P}(n)$ is a decreasing function of n [121]. The variance of this distribution is;

$$\mathrm{Var}\,(n) = \langle n \rangle^2 + \langle n \rangle \tag{D.14}$$

In comparing the variance (equation D.14) with that of the coherent state (Eqn D.6), it is obvious that this incoherent source is inherently more 'noisy'. Such a state is referred to as super-poissonian. Substitution of equation D.14 into equation D.2 gives the normalized second-order auto-correlation function as;

$$g^{(2)} = 2 \tag{D.15}$$

Although we have considered a specific type of field, namely a mixture of many coherent fields with random phases, this geometric distribution actually describes a much more general set of states. A general 'incoherent' state can be defined as being the most disordered state possible, and is found by maximizing the entropy. The entropy of a state defined by density matrix ρ is given by $S = -\mathrm{Tr}\,(\rho \ln \rho)$. Solving for ρ subject to the constraints $\mathrm{Tr}\rho = 1$ and $\mathrm{Tr}\left(\rho \hat{a}^\dagger \hat{a}\right) = \langle n \rangle$ yields an equation for ρ identical to equation D.12 [139].

Given the plethora of physical systems in which incoherent emission can be generated, it is not surprising that there are many labels for light exhibiting statistics described by equation D.13. Perhaps the most popular derivation of incoherent field statistics is based on a single-frequency mode of black-body radiation. This is referred to as *thermal light*. Alternatively, emission from a collection of independent atoms is known as *chaotic light*. Whatever the physical process generating the emission, the photon statistics are identically described by the geometric distribution.

D.3.1 Photon-bunching

The value of $g^{(2)} = 2$ for incoherent light may seem strange. This suggests that the individual photons are more likely to arrive in pairs than individually, an effect referred to as *photon bunching*. This is somewhat counter-intuitive, since the sources of the emission (e.g. an ensemble of atoms, or an ensemble of laser fields as in equation D.9) are completely independent. Indeed in the early days of quantum optics, the prediction of super-poissonian statistics was a much contended issue (for discussion of this see [139]).

This effect can be understood physically by thinking in terms of field *amplitude* rather than intensity (or number of photons). Recall that the net field amplitude, the complex parameter α, randomly fluctuates depending on the phases of all the constituent fields. If a photon is detected at some instant in time, then it is likely that α has momentarily fluctuated to some large value. This in turn implies the probability for detecting a second photon will be larger than average at that instant, and for some time after the initial photon detection.

D.3.2 Photon-bunching in multimode fields

The super-poissonian statistics described by equations D.14 applies to a single frequency mode of the field. If this is expanded to include N_m modes, then the variance becomes (see §13.3 of [28]);

$$\mathrm{Var}\,(n) = \frac{\langle \hat{n} \rangle^2}{N_m} + \langle \hat{n} \rangle \tag{D.16}$$

Such that if there are many modes of the field simultaneously detected, then the statistics become Poissonian. Again this can be understood in terms of the complex parameter α. Mixing of multiple modes is equivalent to taking the average over multiple *independent* values of α. By increasing the number of modes, the fluctuations in any individual mode will have decreased overall effect.

D.3.3 Dynamics of photon-bunching

For two photon-detection events of incoherent light that are separated in time by t, we expect super-poissonian statistics for short t, however for long t the detection events become statistically independent. This time dependence of the intensity correlations is related to the temporal width of the individual photon wave-packets, and therefore the nature of the source.

The transition from bunching at zero time-delay, where $g^{(2)} = 2$, to a large time when statistics should again reduce to being Poissonian with $g^{(2)}(t) = 1$ is derived in §3.7 of [98]. For light with a gaussian spectral profile, for example emission from a Doppler-broadened spectral line, the time dependence is;

$$g^{(2)}(t) = 1 + \exp\left(\frac{-\pi t^2}{t_c^2}\right) \tag{D.17}$$

And for light with a Lorentizan spectral profile, for example emission from a homogeneously-broadened spectral line, the second-order correlation function is;

$$g^{(2)}(t) = 1 + \exp\left(\frac{-2|t|}{t_c}\right) \tag{D.18}$$

where the rate at which the statistics transition from being super-Poissonian to Poissonian are parameterized by the coherence time t_c. Figure D.2 shows $g^{(2)}$ as a function of time t for incoherent light with gaussian or lorentzian spectra, as well as coherent light.

Another point to consider is in the timing characteristics of the detection process itself. The results presented above refer to excitations of the quantized electromagnetic field, without considering the extents to which these properties can actually be measured.

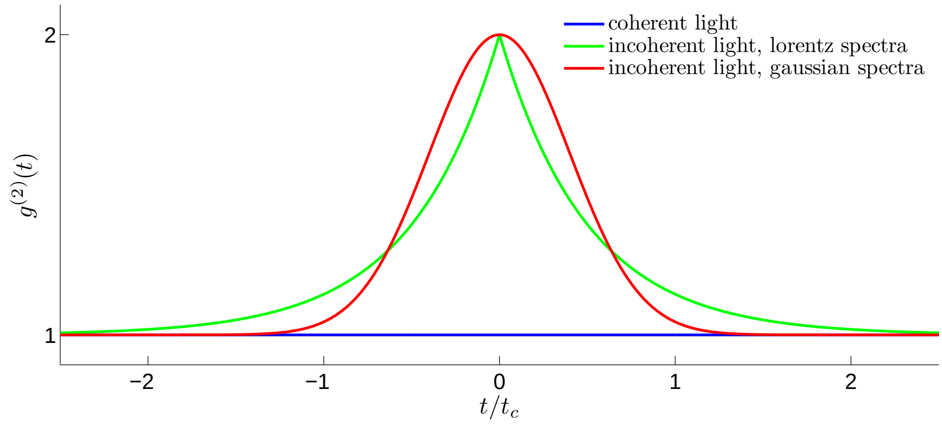

Figure D.2: The second-order correlation function, $g^{(2)}$, for incoherent fields shows photon bunching at time $t = 0$. However the statistics become poissonian when the time between photon detection events is large, because the fields are independent.

If the detector integration-time is denoted as T, and the mean number of detected events during this time is $\langle m \rangle$, then the variance is given by (see §3.9 of [98]);

$$\mathrm{Var}\,(m) = \langle m \rangle + \langle m \rangle^2 \frac{t_c^2}{2T^2} \left(e^{-2T/t_c} - 1 + \frac{2T}{t_c} \right) \tag{D.19}$$

This reduces to equation D.14 in the limit of 'instantaneous' detection where $T \ll t_c$. For $T \gg t_c$, then $\mathrm{Var}\,(m) \approx \langle m \rangle + \langle m \rangle^2 (t_c/T)$. In this limit any bunching effects will be averaged out, and the statistics approach the Poissonian distribution.

D.4 Number state

A number, or Fock state, is a state of the field with a well-defined photon number. Such a state is denoted as $|n\rangle$, and has a mean photon number of n. Since the photon number is exactly known, the variance of this distribution is zero.

$$\langle \hat{n} \rangle = n \tag{D.20}$$

$$\mathrm{Var}\,(\hat{n}) = 0 \tag{D.21}$$

Although Fock states are easy to comprehend, they represent more of a holy-grail than a physical reality in quantum optics. As discussed in the introductory chapter (see §1.2), there are a number of ways to experimentally generate approximate number-states.

From equations D.21 and D.2, the second-order correlation function is immediately found to be;

$$g^{(2)}(0) = 1 - \frac{1}{n} \tag{D.22}$$

As discussed in §D.2, for a classical model of the field, the lower bound for $g^{(2)}(0)$ is 1, which is assuming Poissonian statistics. Thus a value of $g^{(2)}(0) < 1$ represents a state that is not possible in a classical picture, and is distinctly quantum-mechanical in nature. Such a state is said to exhibit 'sub-poissonian' statistics.

APPENDIX **E**

From detector clicks to correlation functions

This appendix describes the algorithms applied to calculate second-order correlation functions $g^{(2)}$ given a recorded log of photon counts detected by the APD detectors. The initial sections outline how $g^{(2)}$ auto- and cross-correlation functions for a single time difference between two fields can be calculated with only one detector.

These methods are applied in Chapter 7 (specifically, in §7.7) to determine correlations between ASE and RASE modes as well as the auto-correlation functions for each mode individually.

E.1 Auto-correlations with a single detector

In the RASE-experiment setup described in Chapter 7 there is a single APD per field-mode. Ordinarily, measuring the auto-correlations of the individual fields in similar experiments (for example, DLCZ experiments [62]) is performed using a beam-splitter and two APDs per mode arranged as depicted in figure E.1. However the correlation times expected in this experiment ($\sim \mu$s) are much longer than the resolution time of the detection system (20 ns). During post-processing the photon arrival times are sorted into time-bins on the order of μs. In this arrangement, particularly with low count rates, a single APD can be considered as a number-resolving detector through temporal-multiplexing.

In a given time-bin with average photon number n, the auto-correlation can be calculated by *simulating* the counts that would be detected by two APDs if this photon stream was incident on a 50:50 beam-splitter as in figure E.1. The statistics of random partitioning of photons at a beam-splitter is given by the binomial distribution. With exactly n photons incident on an ideal (lossless) 50:50 beam-splitter, the probability of having m_T photons transmitted is (see for example [121]);

$$\mathcal{P}(m_T|n) = \frac{n!}{m_T!(n-m_T)!}\frac{1}{2^n}, \qquad m_T = 0, 1, \ldots, n \tag{E.1}$$

In which case the number of photons reflected is $m_R = n - m_T$. The mean number of photons at each output is related to the mean of the input distribution as;

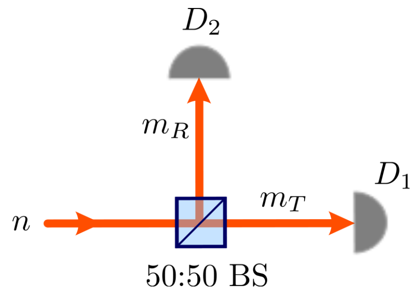

Figure E.1: Second-order auto-correlation functions are determined from the counts measured using a single APD by simulating that this photon stream is incident on a 50:50 beam-splitter and counted by two APDs at the output ports. With n photons incident, the probability of getting m_T photons transmitted and $m_R = n - m_T$ reflected is given by the binomial distribution.

$$\langle m_T \rangle = \langle m_R \rangle = \frac{\langle n \rangle}{2} \tag{E.2}$$

Then the second order correlation can be calculated by summing over all possible values of m_T and m_R weighted by the probability distribution both of the input field, $\mathcal{P}_0(n)$, and the beam-splitter partitioning.

$$g^{(2)} = \frac{\langle m_T m_R \rangle}{\langle m_T \rangle \langle m_R \rangle} \tag{E.3}$$

$$= \frac{4}{\langle n \rangle^2} \sum_{n=2}^{\infty} \mathcal{P}_0(n) \sum_{m_T=1}^{n-1} \mathcal{P}(m_T|n) m_T \left(n - m_T\right) \tag{E.4}$$

Here the limits on the summations are chosen as the first and last possible values that give a non-zero result. For instance, if $n = 1$, then there can never be a 'coincidence' event since (m_T,m_R)=(0,1) or (1,0) and in both cases $\langle m_T m_R \rangle = 0$. Rewriting equation E.4 explicitly in terms of experimental count values and expanding the summation terms gives;

$$g^{(2)} = \frac{4N_{shots}}{N_{total}^2} \left\{ \frac{1}{2}N_2 + \frac{3}{2}N_3 + 3N_4 + 5N_5 + \ldots \right\} \tag{E.5}$$

where N_{shots} is the total number of trials in the experiment, and N_{total} is the sum of all counts, such that $\langle n \rangle = N_{total}/N_{shots}$. The number of instances when there was x counts in one time-bin is denoted as N_x, and the probability distribution of n is determined from these measurements as $\mathcal{P}_0(x) = N_x/N_{shots}$.

E.2 Cross-correlations

Calculating correlation functions between different field modes, separated in space or time or both, is obtained simply from the direct multiplication of the measured counts as;

$$g^{(2)} = \frac{\langle n_a n_b \rangle}{\langle n_a \rangle \langle n_b \rangle} \tag{E.6}$$

In general, it is useful to average not only across many experiment trials, but also across the time domain within a single trial itself. In this case, the correlation function is re-defined as;

$$g^{(2)}(\tau) = \begin{cases} \sum_{t_0}^{t_{max}} \frac{\langle n_A(t) n_R(t+\tau) \rangle}{\langle n_A(t) \rangle \langle n_R(t+\tau) \rangle} & \tau \geq 0 \\ \sum_{t_0}^{t_{max}} \frac{\langle n_A(t-\tau) n_R(t) \rangle}{\langle n_A(t-\tau) \rangle \langle n_R(t) \rangle} & \tau < 0 \end{cases} \tag{E.7}$$

where the subscripts A and R represent ASE and RASE modes respectively, and t is the time relative to the rephasing π-pulses. The parameter τ represents a shift or lag between the two modes. The bounds on the sum are limited by the finite length of the signal recorded in a single trial T, and the lag time τ as $t_{max} \leq T - |\tau|$. Note that the different forms given in equation E.7 which depend on the sign of τ would be formally equivalent if we had infinitely long signals. For clarity, throughout the remainder of this section the functional forms will be written assuming positive τ.

The normalization of $g^{(2)}$ up to this point has been written assuming complete independence between trails. In reality, there are slowly varying experimental factors, such as fluctuating laser intensity, which can result in a higher degree of correlation between neighbouring trials than for trials that are further separated. To partially account for this, the correlation functions can be normalized to a cross-correlation between signals in adjacent trials.

$$g^{(2)}(\tau) = \sum_{t_0}^{t_{max}} \frac{\left\langle n_A^i(t) n_R^i(t+\tau) \right\rangle}{\left\langle n_A^i(t) n_R^{i+1}(t+\tau) \right\rangle} \tag{E.8}$$

where i indexes the trial. Furthermore, to reduce the additional statistical noise that will be introduced by normalizing this way, the denominator can be instead calculated as an average correlation between a small number of neighbouring trials M [62].

$$g^{(2)}(\tau) = \sum_{t_0}^{t_{max}} \frac{\left\langle n_A^i(t) n_R^i(t+\tau) \right\rangle}{\frac{1}{M} \sum_{m=1}^{M} \left\langle n_A^i(t) n_R^{i+m}(t+\tau) \right\rangle} \tag{E.9}$$

For the data presented in this thesis, M was typically on the order of 10. This equation is also valid for finding auto-correlations as a function of lag time τ, except that for the point $\tau = 0$, the simulated beam-splitter measurement described by equation E.5 is used.

APPENDIX F

Summary of laboratory equipment

This appendix summarizes the important aspects of the laboratory setup used for the experiments described throughout this thesis.

F.1 Optics

Schematic

The layout of the optical components for the RASE experiment (described in Chapter 7) is shown in figure F.1. The detection system is interchangeable between heterodyne or single-photon-counting. In essence, there are 6 beam paths with independently controllable frequency. Of these, 5 go to the Pr:YSO sample, and one is a local oscillator for the heterodyne detection.

The acronyms used in figure F.1 have the following meanings;

BS	50:50 beam-splitter
PBS	polarizing beam-splitter
HWP	half waveplate
QWP	quarter waveplate

Laser

The laser is a Coherent 699 ring dye laser, which has been modified to have a sub-kilohertz linewidth [141].

Mechanical shutter

A shutter was necessary for ensuring the optical fibre connection into the APD is not exposed to any bright ($\sim$mW) laser beams during the experiment. If this was not the case, then long-lived fluorescence was observed with the APD detectors, thought to originate from the glue in the fibre connector.

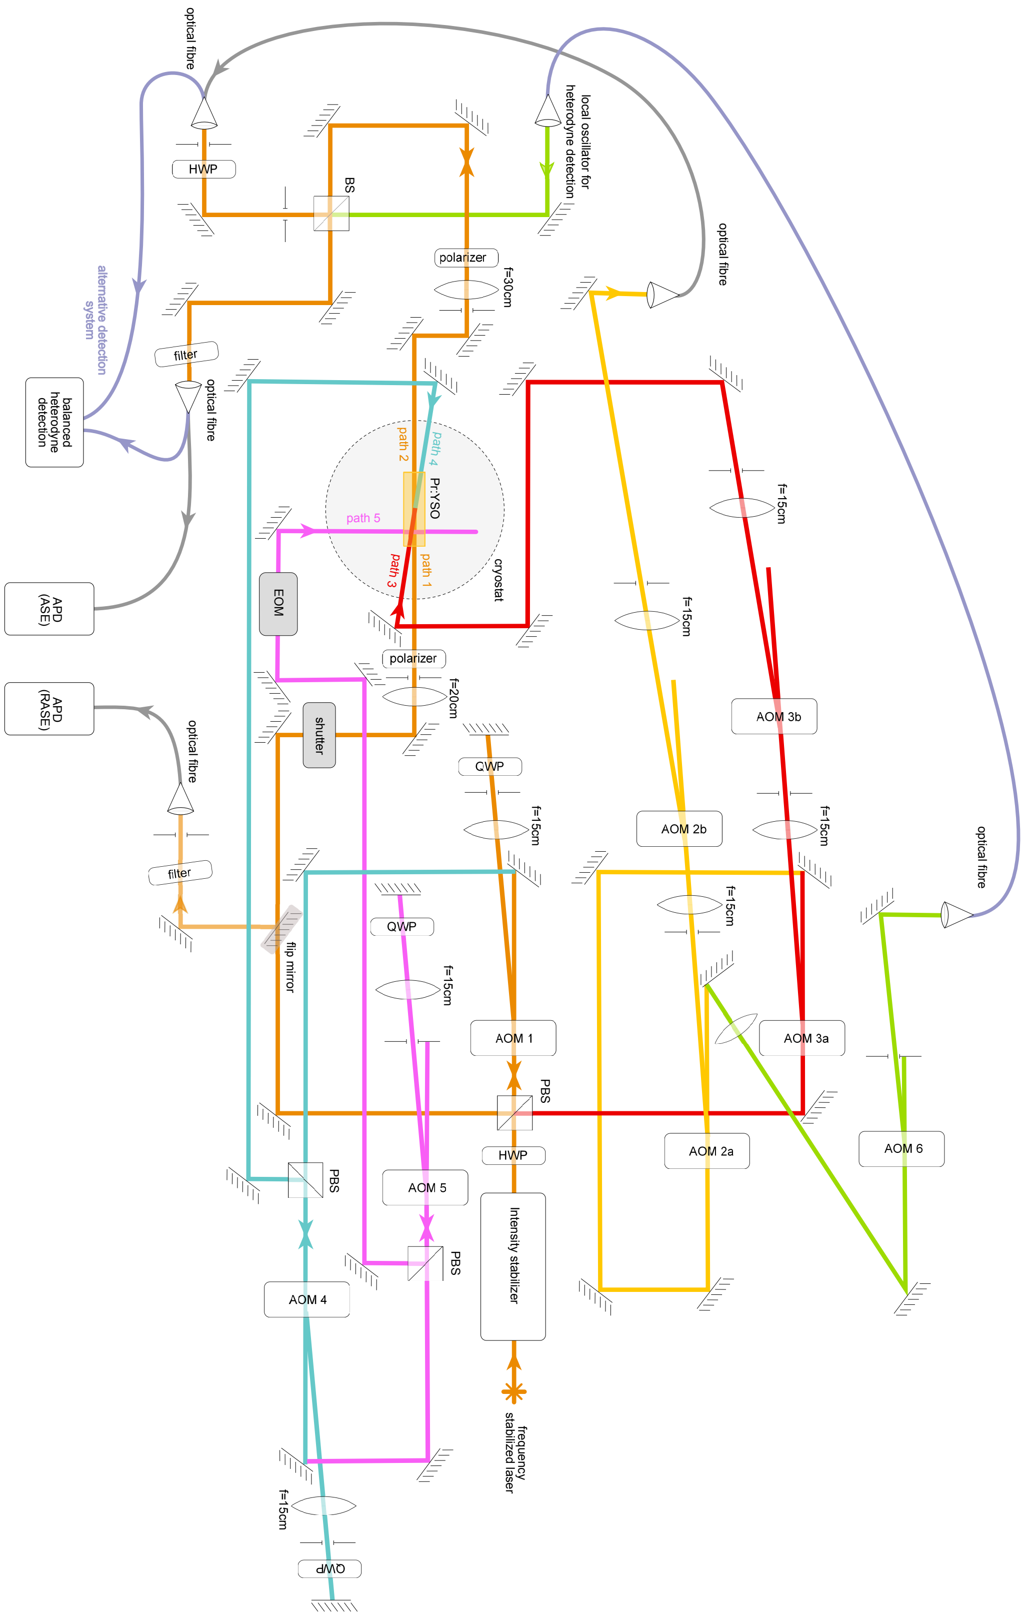

Figure F.1: Schematic of optics experiment for counter-propagating RASE experiment described in Chapter 7. All other experiments described in this thesis were performed with subsets of this setup.

A UniBlitz shutter with 2 mm aperture was used for this purpose. This shutter is mechanical, with a maximum frequency rating of 100 Hz (for continuous operation), and a opening/closing speed of 300 µs.

Bandpass filters

Direct transitions from $^3H_4 \rightarrow {}^1D_2$ in Pr^{3+}, at a wavelength of 606 nm, occur with a probability of only ∼3%. The majority of ions will relax to the ground state via intermediate crystal field levels. A measured spectra of 0.005% Pr^{3+} : Y_2SiO_5 is shown in figure F.3. The sample was excited at 606 nm, while the laser was scanning over the edge of the 5 GHz wide inhomogeneous line (to minimize re-absorption of the emission at 606 nm) and the emission was collected from a perpendicular direction to the pump beam through a fibre to a CCD spectrometer.

The photon-counting detectors provide no frequency discrimination in this part of the spectrum (see §F.3), therefore a filter is required to transmit light at 606 nm while attenuating from approximately 610 nm onwards. The filter used was a 10 nm bandpass filter from Thorlabs; FB610-10. This filter consists of a thin film Fabry-Perot-like interferometer (to achieve the narrow transmission window) and an additional blocking component (to attenuate further away radiation). The transmission window can be shifted in frequency by tilting the filter. Therefore for the experiments described in this thesis, the filter was tilted to shift the long-wavelength edge of the transmission window to 606 nm, such that this emission would be maximally transmitted, while the emission at longer wavelengths would be heavily attenuated. Although the specified maximum transmission is ≈55%, the measured value was closer to 60% at the angle optimized for 606 nm.

Acousto-optic modulators (AOMs)

AOMs were used in all experiments described in this thesis for modulating the intensity and frequency of the optical beams. In most cases, the specific model used was the IntraAction ATM-80A1. The Acousto-optic material in this AOM is Tellurium Dioxide. The center frequency is 80 MHz, and depending on the experiment the AOMs were operated in the 60-90 MHz range. The peak diffraction efficiency is 85%. Typically, the delay between switching on the RF and the optical beam being diffracted was ∼µs.

As shown in figure F.1, in some cases a double-pass AOM configuration was used. Such an arrangement is used to minimize the relative difference in beam pointing that occurs when the AOM is driven at different frequencies.

Electro-optic phase modulator (EOM)

In the experiment described in Chapter 7, an EOM was used to sweep the frequency of the filter-burning field (path 5 in figure F.1).

This EOM was a Thorlabs model, operating at 400-600 nm. The driving RF for the

EOM was sourced from a Stanford DS345 30MHz synthesized function generator, amplified by a DC-5 MHz high voltage (±150 V) amplifier from Falco Systems.

F.2 Optical amplitude and frequency control

RF generators

There were various RF sources used in the experiments, each having different functionality. The three sources most often used were;

- Home-made direct digital synthesis (DDS) systems, labelled as 'J850'.
  Each channel of the J850 source can output a single frequency tone, or can alternate/sweep between two frequency tones. The sweeping functionality is particularly useful for hole-burning and state preparation. The clock frequency is 300 MHz, generated from a 15 MHz external reference of < 0.01 ppm stability (A Stanford function generator, model DS345). A maximum of 6 J850 channels were used simultaneously.

- Spincore RadioProcessor.
  This has a single channel DDS output with 14-bit resolution. The clock frequency is also 300 MHz. Up to 16 frequencies and phases can be programmed, as well as a pulse sequence for switching the RF output on/off or switching between different frequencies. The RF output can also be amplitude shaped. This RF source was used to generate echo sequences, or probe-pulse sequences, usually with gaussian shaped amplitude profiles. The internal clock was upgraded to an oven-controlled crystal-oscillator stable to ±0.2 ppm.

- Tektronix arbitrary waveform generator (AWG).
  This single-channel AWG has a clock frequency of 1 GHz, and a stability of ±1 ppm. There are two output signals, with one being the inverse of the other. This AWG was used similarly to the RadioProcessor for generating sequences of shaped pulses.

In order to combine and switch the signals from the various sources to the different AOMs, a network of switches (Mini-Circuits high isolation TTL switch ZASWA-2-50DR+) were used. The RF signal is passed through a low-pass filter (Mini-Circuits BLP 100 or 90 MHz) and amplified (Mini-Circuits 29.5 dBm ZHL-3A+ amplifier) before being sent to the AOMs.

The central timing unit used to switch the RF sources, and to trigger the RadioProcessor and/or AWG was a Spincore PulseBlaster DDS (usb). This has 10 TTL outputs, and a 100 MHz clock frequency.

F.3 Detection

Single-photon counting modules

All photon-counting measurements described in this thesis were performed with Perkin Elmer Avalanche Photodiode (APD) detectors, model SPCM-AQR-14-FC. For this model the manual specifies an afterpulsing probability of 0.5%. The individual specifications for the two APDs labelled in this thesis as D_A and D_R (see for example Chapter 7) as measured by Perkin Elmer are;

	D_A	D_R
connectorized photon detection efficiency (at 550 nm)	65.4 %	60.4 %
dark counts	54 Hz	83 Hz
dead time	49.9 ns	48.6 ns

According to the generalized efficiency Vs wavelength graph in the manual, the efficiency at 606 nm should increase to $\sim$67%, however the fibre connector lens is optimized for 550 nm, so the detection efficiency will be slightly reduced.

Other photo-detectors

Heterodyne detection involves spatially overlapping the signal field to be measured with a bright 'local-oscillator' beam, and then the signal amplitude and phase is given by the properties of the radio-frequency beat between the two light fields.

The local oscillator was combined with the signal beam using a 50:50 beam-splitter. Then depending on the experimental arrangement, either a single output port of the beam-splitter was incident on a photo-detector, or alternatively the two outputs were detected and the photocurrents subtracted for so called *balanced* heterodyne detection, in which classical noise on the local oscillator is subtracted out.

The two detectors and some of their specifications are listed below;

- For balanced heterodyne detection:
  Thorlabs PDB120A balanced amplified photo-detector
  320-1000 nm wavelength range
  DC - 75 MHz bandwidth
  Transimpedance gain 180 V/mA
  Common mode rejection ratio > 35 dB

- For single-mode heterodyne detection:
  New Focus low noise photo-receiver model 1801, AC-coupled
  320-1000 nm wavelength range
  25 kHz - 125 MHz bandwidth
  Transimpedance gain 40 V/mA

F.4 Data acquisition

Analog signal recording

When heterodyne detection is being used, the data to be recorded is an analog voltage signal output from the photo-detectors. A digital USB oscilloscope was used for this purpose. The *cleverscope* CS328A has the following specifications;

- 100 MHz sampling rate
- 2 × 14 bit analog inputs
- 8 × digital inputs
- 4 Mega samples of storage per channel
- realtime FFT provides spectral analysis of signal

Although the cleverscope is a very affordable unit, and is extremely useful as an oscilloscope and spectrum analyser, it is problematic to implement for continuously recording data for lengthy acquisition periods. Firstly, there is no capability in the software to automate acquisitions that require many cycles of filling the 8 Mega samples on-board buffer. Also, the full 100 MHz sampling rate cannot be accessed for time-frames greater than 50 μs. Most of the heterodyne data presented in this thesis was for a single-shot time window of 100 μs, where the maximum sampling rate drops to 50 MHz.

To summarize, the cleverscope is a great oscilloscope and spectrum analyser, but is not an ideal unit for data acquisition.

Digital signal recording

For experiments where single-photon-counting detection is being used, the data to be recorded is the arrival time of the digital (TTL) output pulses from the APD detectors. The data acquisition in this case was implemented using a Xylo-EM FPGA development board, with open source firmware (verilog) and software (c/c++) written by Polyakov *et al.* [142], and a GUI written in python.

The FPGA is configured to have one counter-reset (start) channel, and four time-stamping (stop) channels. When a TTL pulse is input onto any of the four stop channels, the time relative to the last start TTL is recorded, in increments of 20.8 ns (using the 48 MHz internal clock of the FPGA). This board provides an inexpensive, versatile, and easy-to-implement data acquisition system for timestamping pulse arrival times from the APD detectors.

In post-processing the large sets of data recorded for the RASE experiment described in Chapter 7, it was seen that there was a higher probability than expected of recording a count immediately following another count. These spurious counts occur on time scales longer than can be attributed to detector afterpulsing, where a false count is registered immediately after the detector deadtime. Instead these counts and are thought to be due to electrical ringing effects, probably due to the use of 2 m long BNC cables for connecting the APD signal output to the FPGA stop-channel inputs.

To benchmark the performance of the timestamping detection system in an arrangement similar to the experiment setup, a low level of multimode light (count rate of 100 Hz) was incident on the APDs, and recorded by the FPGA via 2 m long cables terminated with a 50 Ω connection to ground in parallel at each end. The resulting probability of detecting a count within a certain delay of an initial count is shown in figure F.4. For detector D_A, there is a large probability of getting a count in the FPGA clock cycle immediately following a prior count, which is within the detector deadtime and therefore cannot be an afterpulse. Also for both detector channels, the probability remains higher than the level expected for actual photon detection events for approximately 200 ns.

This high count probability lasting out to 200 ns is the reason for the artificially high auto-correlations measured for the RASE experiment (see figure 7.17), since this calculation is based on the number of 2-photon events occurring in the same detector in a 1 μs timebin. Before applying this detection/data-acquisition system for determining auto-correlations, this problem of false counts on timescales longer than the detector deadtime needs to be further investigated. A potential fix would be to use shorter cables to input the TTLs from the detector into the FPGA, or have intermediate line-drivers.

F.5 Cryostat

The cryostat used for the majority of experiments presented in this thesis was a Janis bath cryostat, and in most cases the sample was fully immersed in liquid helium. The helium in the sample space is evaporatively cooled by pumping with a mechanical vacuum pump, such that it becomes superfluid. The cryostat has a 6 litre capacity for liquid helium, and 7 litres for liquid nitrogen. Under typical operating conditions, the helium hold time of the cryostat was approximately 30 hours.

Pressure fluctuations transferred onto the sample via the liquid helium cause instability in measurements, particularly for coherent transient signals. Maximum stability was achieved by closing the needle valve between the helium reservoir and the sample space while sensitive measurements were being made. This included the many hours of data recorded for the RASE experiment described in Chapter 7.

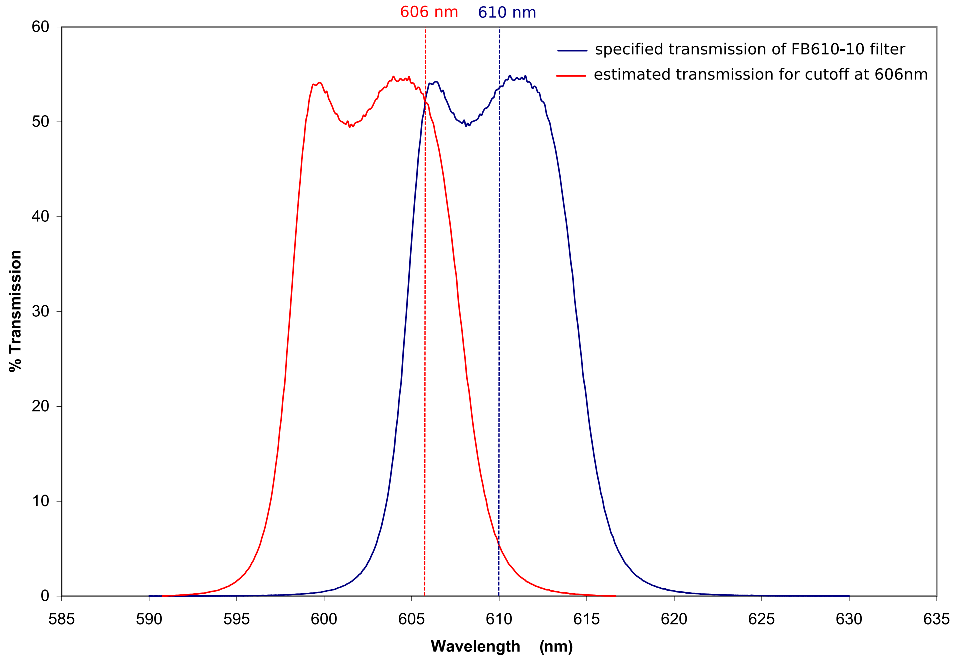

Figure F.2: Transmission curve for thorlabs bandpass filter FB610-10, where at normal incidence the transmission window is centred at 610 nm and has a fwhm of 10 nm. The red curve is the same profile, but shifted to approximate the spectral transmission window when the filter is tilted to optimize transmission at 606 nm, while providing good extinction at longer wavelengths.

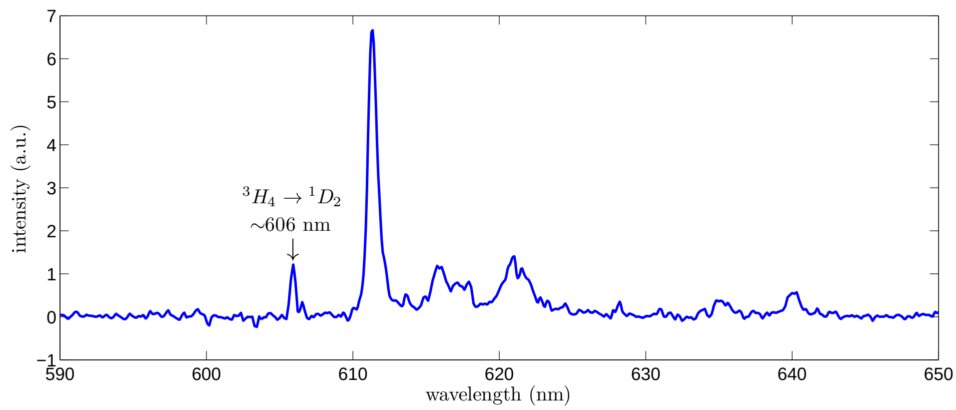

Figure F.3: Pr^{3+} : Y_2SiO_5 spectrum measured using a CCD spectrometer. The sample was excited through the front window of the cryostat, with the laser scanning across a 6 GHz range towards the edge of the $^3H_4 \rightarrow {}^1D_2$ transition, and the spectrometer detecting emission through the side window of the cryostat. The process was repeated with the laser locked, and therefore hole-burning occurring such that there is no emission, and this background spectrum was subtracted from the emission spectrum to give the above plot. The direct transition from $^3H_4 \rightarrow {}^1D_2$ is visible, and the majority of the emission occurs at wavelengths 612 nm and longer.

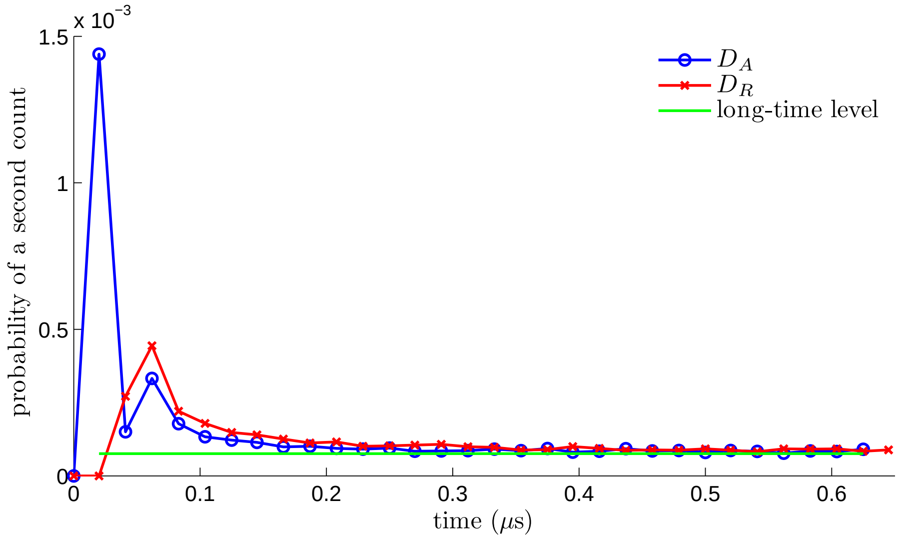

Figure F.4: Probability as a function of delay of detecting a second count following an initial count for each of the APDs for multimode incident light, recorded using the FPGA. The discrete points in time represent each clock tick of the FPGA, ≈20 ns apart. There is a higher chance of getting a count in the 200 ns period following an initial count (slightly more so in D_R), and there is a large spike in the time-bin immediately following a prior count for detector D_A. This is attributed to electronic noise. The green line shows the average probability level between 50 and 100 μs following a pulse, which is the same for both detectors.

Bibliography

[1] S. E. Beavan, E. Fraval, M. J. Sellars and J. J. Longdell, "Demonstration of the reduction of decoherent errors in a solid-state qubit using dynamic decoupling techniques" *Physical Review A* **80**, 032308, (2009).
(Cited on pages xvii, 22, 30, 167 and 168)

[2] P. M. Ledingham, W. R. Naylor, J. J. Longdell, S. E. Beavan and M. J. Sellars, "Nonclassical photon streams using rephased amplified spontaneous emission" *Physical Review A* **81**, 012301, (2009).
(Cited on pages xvii, 10, 19 and 83)

[3] S. E. Beavan, P. M. Ledingham, J. J. Longdell and M. J. Sellars, "Photon echo without a free induction decay in a double-Λ system" *Optics Letters* **36**, 1272–4, (2011).
(Cited on page xvii)

[4] D. L. McAuslan, P. M. Ledingham, W. R. Naylor, S. E. Beavan, M. P. Hedges, M. J. Sellars and J. J. Longdell, "Photon-echo quantum memories in inhomogeneously broadened two-level atoms" *Physical Review A* **84**, 022309, (2011).
(Cited on pages xvii, 32 and 61)

[5] D. L. McAuslan, *Quantum Computing Hardware Based on Rare-Earth-Ion Doped Whispering-Gallery Mode Resonators*, Ph.D. thesis, The University of Otago, (2011).
(Cited on page xvii)

[6] P. M. Ledingham, *Optical Rephasing Techniques with Rare Earth Ion Dopants for Applications in Quantum Information Science*, Ph.D. thesis, The University of Otago, (2011).
(Cited on pages xvii, 19, 20, 59, 81, 82 and 175)

[7] A. Einstein, B. Podolsky and N. Rosen, "Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?" *Physical Review* **47**(10), 777–780, (1935).
(Cited on page 1)

[8] J. S. Bell, "On the Problem of Hidden Variables in Quantum Mechanics" *Reviews of Modern Physics* **38**(3), 447–452, (1966).
(Cited on page 1)

[9] P. W. Shor, "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer" *SIAM Journal on Computing* **26**, 1484–1509, (1997).
(Cited on page 1)

[10] R. P. Feynman, "Simulating physics with computers" *International Journal of Theoretical Physics* **21**, 467–488, (1982).
(Cited on page 1)

[11] T. Monz, P. Schindler, J. Barreiro, M. Chwalla, D. Nigg, W. Coish, M. Harlander, W. Hänsel, M. Hennrich and R. Blatt, "14-Qubit Entanglement: Creation and Coherence" *Physical Review Letters* **106**(13), 1–4, (2011).
(Cited on page 1)

[12] R. L. Rivest, A. Shamir and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems" *Commun. ACM* **21**(2), (1978).
(Cited on page 1)

[13] S. Wiesner, "Conjugate coding" *SIGACT News* **15**(1), 78–88, (1983).
(Cited on page 2)

[14] C. H. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing"in "Proceedings of IEEE International Conference on Computers Systems and Signal Processing," pp. 175–179, Bangalore, India, (1984).
(Cited on page 2)

[15] A. Muller, J. Breguet and N. Gisin, "Experimental demonstration of quantum cryptography using polarized photons in optical fibre over more than 1 km" *Europhysics Letters* **23**, (1993).
(Cited on page 2)

[16] D. Bouwmeester, A. Ekert and A. Zeilinger, *The Physics of Quantum Information: Quantum Cryptography, Quantum Teleportation, Quantum Computation*, Springer-Verlag, (2001).
(Cited on pages 2, 3 and 11)

[17] A. Ekert, "Quantum cryptography based on Bell's theorem" *Physical Review Letters* **67**(6) 661–663, (1991).
(Cited on page 2)

[18] A. Ekert, J. Rarity, P. Tapster and G. Massimo Palma, "Practical quantum cryptography based on two-photon interferometry" *Physical Review Letters* **69**(9), 1293–1295, (1992).
(Cited on page 2)

[19] N. Gisin, G. Ribordy, W. Tittel and H. Zbinden, “Quantum cryptography” *Reviews of Modern Physics* **74**(1), 145–195, (2002).
(Cited on pages 2 and 6)

[20] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, Cambridge, (2000).
(Cited on pages 2 and 168)

[21] F. Grosshans and P. Grangier, “Continuous Variable Quantum Cryptography Using Coherent States” *Physical Review Letters* **88**(5), 1–4, (2002).
(Cited on page 3)

[22] G. Leuchs, N. J. Cerf and E. Polzik, *Quantum information with continuous variables of atoms and light*, Imperial College Press, London, (2007).
(Cited on page 3)

[23] H. Lo, X. Ma and K. Chen, “Decoy State Quantum Key Distribution” *Physical Review Letters* **94**(23), 15–18, (2005).
(Cited on page 3)

[24] W. K. Wootters and W. H. Zurek, “A single quantum cannot be cloned” *Nature* **299**, 802–803, (1982).
(Cited on page 3)

[25] D. Dieks, “Communication by EPR devices” *Physics Letters A* **92**(6), 271–272, (1982).
(Cited on page 3)

[26] A. Zeilinger, G. Weihs, T. Jennewein and M. Aspelmeyer, “Happy centenary, photon.” *Nature* **433**, 230–8, (2005).
(Cited on page 3)

[27] M. Fox, *Quantum Optics: An Introduction*, Oxford University Press, (2006).
(Cited on pages 4 and 202)

[28] L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics*, Cambridge University Press, (1995).
(Cited on pages 4, 42, 51, 65, 125, 200, 202 and 205)

[29] B. Lounis and M. Orrit, “Single-photon sources” *Reports on Progress in Physics* **68**(5), 1129–1179, (2005).
(Cited on pages 4 and 5)

[30] J. G. Rarity, P. C. M. Owens and P. R. Tapster, “Quantum Random-number Generation and Key Sharing” *Journal of Modern Optics* **41**(12), 2435–2444, (1994).
(Cited on page 4)

[31] A. Stefanov, N. Gisin, O. Guinnard, L. Guinnard and H. Zbinden, "Optical quantum random number generator" *Journal of Modern Optics* **47**(4), 595–598, (2000).
(Cited on page 4)

[32] E. Knill, R. Laflamme and G. J. Milburn, "A scheme for efficient quantum computation with linear optics." *Nature* **409**, 46–52, (2001).
(Cited on page 4)

[33] C. Santori, D. Fattal and Y. Yamamoto, *Single-photon Devices and Applications*, John Wiley & Sons, (2010).
(Cited on page 5)

[34] C. Hong and L. Mandel, "Experimental realization of a localized one-photon state" *Physical Review Letters* **56**(1), 58–60, (1986).
(Cited on page 6)

[35] R. W. Boyd, *Nonlinear Optics*, Academic Press, (2003).
(Cited on page 7)

[36] M. Scully and S. Zubairy, *Quantum Optics*, Cambridge University Press, 1st edition, (1997).
(Cited on pages 44, 46, 66, 189 and 190)

[37] G. Walls, D.F. , Milburn, *Quantum optics*, Springer-Verlag, Berlin, (1994).
(Cited on page 7)

[38] Z. H. Levine, J. Fan, J. Chen, A. Ling and A. Migdall, "Heralded, pure-state single-photon source based on a Potassium Titanyl Phosphate waveguide" *Optics Express* **18**(4), 3708–18, (2010).
(Cited on page 7)

[39] D. J. Wineland, W. M. Itano and J. C. Bergquist, "Absorption spectroscopy at the limit: detection of a single atom" *Optics Letters* **12**(6), 389–91, (1987).
(Cited on page 8)

[40] S. Hell and E. H. K. Stelzer, "Properties of a 4Pi confocal fluorescence microscope" *Journal of the Optical Society of America A* **9**(12), 2159, (1992).
(Cited on page 8)

[41] I. Ichimura, S. Hayashi and G. S. Kino, "High-density optical recording using a solid immersion lens." *Applied Optics* **36**(19), 4339–48, (1997).
(Cited on page 8)

[42] S. van Enk, "Atoms, dipole waves, and strongly focused light beams" *Physical Review A* **69**, 043813, (2004).
(Cited on page 8)

[43] M. Sondermann, R. Maiwald, H. Konermann, N. Lindlein, U. Peschel and G. Leuchs, "Design of a mode converter for efficient light-atom coupling in free space" *Applied Physics B* **89**(4), 489–492, (2007).
(Cited on page 8)

[44] R. Maiwald, D. Leibfried, J. Britton, J. C. Bergquist, G. Leuchs and D. J. Wineland, "Stylus ion trap for enhanced access and sensing" *Nature Physics* **5**(8), 551–554, (2009).
(Cited on page 8)

[45] N. Lindlein, R. Maiwald, H. Konermann, M. Sondermann, U. Peschel and G. Leuchs, "A new 4π geometry optimized for focusing on an atom with a dipole-like radiation pattern" *Laser Physics* **17**(7), 927–934, (2007).
(Cited on page 8)

[46] K. Lee, X. Chen, H. Eghlidi, P. Kukura, R. Lettow, A. Renn, V. Sandoghdar and S. Götzinger, "A planar dielectric antenna for directional single-photon emission and near-unity collection efficiency" *Nature Photonics* **5**(3), 166–169, (2011).
(Cited on page 8)

[47] X.-W. Chen, S. Götzinger and V. Sandoghdar, " 99% Efficiency in Collecting Photons From a Single Emitter" *Optics Letters* **36**(18), 3545–7, (2011).
(Cited on page 8)

[48] H. Kimble, "Strong interactions of single atoms and photons in cavity QED" *Physica Scripta* **T76**, 127–137, (1998).
(Cited on page 8)

[49] H. G. Barros, A. Stute, T. E. Northup, C. Russo, P. O. Schmidt and R. Blatt, "Deterministic single-photon source from a single ion" *New Journal of Physics* **11**, 103004, (2009).
(Cited on page 8)

[50] S. M. Spillane, T. J. Kippenberg and K. J. Vahala, "Ultrahigh-Q toroidal microresonators for cavity quantum electrodynamics" *Physical Review A* **71**, 013817, (2005).
(Cited on page 9)

[51] L.-M. Duan, M. D. Lukin, J. I. Cirac and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics" *Nature* **414**, 413–418, (2001).
(Cited on pages 9, 14 and 51)

[52] 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).
(Cited on page 11)

[53] M. P. Hedges, *High Performance Solid State Quantum Memory*, Ph.D. thesis, Australian National University, (2011).
(Cited on pages 11, 28, 31, 32, 34 and 60)

[54] N. Sangouard, C. Simon, H. de Riedmatten and N. Gisin, "Quantum repeaters based on atomic ensembles and linear optics" *Review of Modern Physics* **83**(1), 33–80, (2011).
(Cited on pages 11 and 12)

[55] C. Cabrillo, J. Cirac, P. García-Fernández and P. Zoller, "Creation of entangled states of distant atoms by interference" *Physical Review A* **59**(2), 1025–1033, (1999).
(Cited on page 11)

[56] M. P. Hedges, J. J. Longdell, Y. Li and M. J. Sellars, "Efficient quantum memory for light" *Nature* **465**, 1052–1056, (2010).
(Cited on pages 14, 31, 34, 155 and 159)

[57] M. Hosseini, B. M. Sparkes, G. Hetet, J. J. Longdell, P. K. Lam and B. C. Buchler, "Coherent optical pulse sequencer for quantum applications" *Nature* **461**, 241–245, (2009).
(Cited on page 14)

[58] H. de Riedmatten, M. Afzelius, M. U. Staudt, C. Simon and N. Gisin, "A solid-state light-matter interface at the single-photon level" *Nature* **456**, 773–777, (2008).
(Cited on page 14)

[59] C. Clausen, I. Usmani, F. Bussières, N. Sangouard, M. Afzelius, H. de Riedmatten and N. Gisin, "Quantum storage of photonic entanglement in a crystal" *Nature* **469**, 508, (2011).
(Cited on page 14)

[60] E. Saglamyurek, N. Sinclair, J. Jin, J. A. Slater, D. Oblak, F. Bussières, M. George, R. Ricken, W. Sohler and W. Tittel, "Broadband waveguide quantum memory for entangled photons" *Nature* **469**, 512, (2011).
(Cited on pages 14 and 23)

[61] C. van der Wal, M. Eisaman, A. André, R. Walsworth, D. Phillips, A. Zibrov and M. Lukin, "Atomic memory for correlated photon states" *Science* **301**, 196, (2003).
(Cited on page 14)

[62] A. Kuzmich, W. P. Bowen, A. D. Boozer, A. Boca, C. W. Chou, L.-M. Duan and H. J. Kimble, "Generation of nonclassical photon pairs for scalable quantum communication with atomic ensembles" *Nature* **423**, 731–734, (2003).
(Cited on pages 14, 209 and 211)

[63] C. Chou, S. Polyakov, A. Kuzmich and H. Kimble, "Single-Photon Generation from Stored Excitation in an Atomic Ensemble" *Physical Review Letters* **92**(21), (2004).
(Cited on page 14)

[64] S. V. Polyakov, C. W. Chou, D. Felinto and H. J. Kimble, "Temporal Dynamics of Photon Pairs Generated by an Atomic Ensemble" *Physical Review Letters* **93**, 263601, (2004).
(Cited on page 14)

[65] W. Jiang, C. Han, P. Xue, L.-M. Duan and G.-C. Guo, "Nonclassical photon pairs generated from a room-temperature atomic ensemble" *Physical Review A* **69**, 043819, (2004).
(Cited on page 14)

[66] M. Eisaman, L. Childress, A. André, F. Massou, A. Zibrov and M. Lukin, "Shaping Quantum Pulses of Light Via Coherent Atomic Memory" *Physical Review Letters* **93**, 233602, (2004).
(Cited on page 14)

[67] D. Felinto, C. W. Chou, H. de Riedmatten, S. V. Polyakov and H. J. Kimble, "Control of decoherence in the generation of photon pairs from atomic ensembles" *Physical Review A* **72**, 053809, (2005).
(Cited on page 14)

[68] A. Black, J. Thompson and V. Vuletić, "On-Demand Superradiant Conversion of Atomic Spin Gratings into Single Photons with High Efficiency" *Physical Review Letters* **95**, 133601, (2005).
(Cited on page 14)

[69] D. Matsukevich, T. Chanelière, M. Bhattacharya, S.-Y. Lan, S. Jenkins, T. Kennedy and a. Kuzmich, "Entanglement of a Photon and a Collective Atomic Excitation" *Physical Review Letters* **95**, 040405, (2005).
(Cited on page 14)

[70] J. Laurat, H. de Riedmatten, D. Felinto, C.-W. Chou, E. W. Schomburg and H. J. Kimble, "Efficient retrieval of a single excitation stored in an atomic ensemble" *Optics Express* **14**, 6912–6918, (2006).
(Cited on pages 14 and 149)

[71] 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**, 828–832, (2005).
(Cited on page 14)

[72] K. S. Choi, A. Goban, S. B. Papp, S. J. van Enk and H. J. Kimble, "Entanglement of spin waves among four quantum memories" *Nature* **468**, 412–416, (2010).
(Cited on page 14)

[73] Y. Dudin, A. Radnaev, R. Zhao, J. Blumoff, T. Kennedy and A. Kuzmich, "Entanglement of Light-Shift Compensated Atomic Spin Waves with Telecom Light" *Physical Review Letters* **105**, 260502, (2010).
(Cited on pages 14 and 15)

[74] R. Zhao, Y. O. Dudin, S. D. Jenkins, C. J. Campbell, D. N. Matsukevich, T. A. B. Kennedy and A. Kuzmich, "Long-lived quantum memory" *Nature Physics* **5**, 100–104, (2009).
(Cited on page 15)

[75] M. D. Eisaman, S. Polyakov, M. Hohensee, J. Fan, P. Hemmer and A. Migdall, "Optimizing the storage and retrieval efficiency of a solid-state quantum memory through tailored state preparation" *Proceedings of SPIE* **6780**, 67800K-1, (2007).
(Cited on page 15)

[76] R. Macfarlane, "High-resolution laser spectroscopy of rare-earth doped insulators: a personal perspective" *Journal of Luminescence* **100**, 1–20, (2002).
(Cited on page 23)

[77] T. W. Mossberg, "Time-domain frequency-selective optical data storage." *Optics Letters* **7**(2), 77–9, (1982).
(Cited on pages 23 and 60)

[78] D. L. McAuslan and J. J. Longdell, "Strong-coupling cavity QED using rare-earth-metal-ion dopants in monolithic resonators: What you can do with a weak oscillator" *Physical Review A* **80**, 062307, (2009).
(Cited on page 23)

[79] G. Liu and B. Jacquier, *Spectroscopic Properties of Rare Earths in Optical Materials*, Springer, Berlin, (2005).
(Cited on pages 24, 27 and 28)

[80] R. W. Equall, R. L. Cone and R. M. Macfarlane, "Homogeneous broadening and hyperfine structure of optical transitions in $Pr^{3+} : Y_2SiO_5$" *Physical Review B* **52**, 3963–3969, (1995).
(Cited on pages 25, 26, 27, 28 and 29)

[81] G. H. Dieke and H. M. Crosswhite, "The Spectra of the Doubly and Triply Ionized Rare Earths" *Applied Optics* **2**(7), 675–686, (1963).
(Cited on page 26)

[82] F. Könz, Y. Sun, C. W. Thiel, R. L. Cone, R. W. Equall, R. L. Hutcheson and R. M. Macfarlane, "Temperature and concentration dependence of optical dephasing, spectral-hole lifetime, and anisotropic absorption in $Eu^{3+}:Y_2SiO_5$" *Physical Review B* **68**, 085109, (2003).
(Cited on page 28)

[83] R. D. Shannon and C. T. Prewitt, "Effective ionic radii in oxides and fluorides" *Acta Crystallographica* **B25**, 925, (1969).
(Cited on page 28)

[84] "Webelements: the periodic table on the web", (accessed December 2011) http://www.webelements.com/silicon/isotopes.html
(Cited on page 28)

[85] R. M. Macfarlane, D. P. Burum and R. M. Shelby, "New Determination of the Nuclear Magnetic Moment of ^{141}Pr" *Physical Review Letters* **49**(9), 636–639, (1982).
(Cited on page 28)

[86] E. Fraval, *Minimising the Decoherence of Rare Earth Ion Solid State Spin Qubits*, Ph.D. thesis, Australian National University, (2006).
(Cited on pages 28, 29, 30 and 169)

[87] J. Bartholomew, personal communication, Australian National University, (2011).
(Cited on page 28)

[88] M. Nilsson, L. Rippe, S. Kröll, R. Klieber and D. Suter, "Hole-burning techniques for isolation and study of individual hyperfine transitions in inhomogeneously broadened solids demonstrated in $Pr^{3+}:Y_2SiO_5$" *Physical Review B* **70**, 214116, (2004).
(Cited on pages 29, 103 and 107)

[89] E. Fraval, M. J. Sellars and J. J. Longdell, "Method of Extending Hyperfine Coherence Times in $Pr^{3+}:Y_2SiO_5$" *Physical Review Letters* **92**, 077601, (2004).
(Cited on pages 30 and 169)

[90] R. Ahlefeldt, *Ion-ion interactions in $EuCl_3 \cdot 6H_2O$*, Honours thesis, Australian National University, (2008).
(Cited on page 32)

[91] F. Graf, A. Renn, U. Wild and M. Mitsunaga, "Site interference in Stark-modulated photon echoes" *Physical Review B* **55**(17), 11 225, (1997).
(Cited on pages 32, 34, 130 and 157)

[92] C. Cohen-Tannoudji, B. Diu and F. Laloe, *Quantum Mechanics*, Wiley-Interscience, (1978).
(Cited on page 42)

[93] Y. Yamamoto and A. Imamoglu, *Mesoscopic quantum optics*, Wiley, (1999).
(Cited on pages 44, 65 and 71)

[94] J. H. Eberly and L. Allen, *Optical resonance and two-level atoms*, Dover Publications, (1975).
(Cited on pages 44, 46, 51, 59, 125 and 194)

[95] C. Cohen-Tannoudji, J. Dupont-Roc and G. Grynberg, *Atom-Photon Interactions*, Wiley-Interscience, (1998).
(Cited on page 46)

[96] D. Griffiths, *Introduction to Electrodynamics*, 3rd edition, Prentice-Hall, (1999).
(Cited on pages 46 and 47)

[97] N. E. Rehler and J. H. Eberly, "Superradiance" *Physical Review A* **3**, 1735–1751, (1971).
(Cited on pages 51 and 194)

[98] R. Loudon, *The quantum theory of light*, 3rd edition, Oxford University Press, (2000).
(Cited on pages 51, 86, 205 and 206)

[99] E. L. Hahn, "Spin Echoes" *Physics Today* **3**(12), 21, (1950).
(Cited on pages 53 and 57)

[100] N. A. Kurnit, I. D. Abella and S. R. Hartmann, "Observation of a Photon Echo" *Physical Review Letters* **13**(19), 567–568, (1964).
(Cited on pages 53, 58 and 59)

[101] I. D. Abella, N. A. Kurnit and S. R. Hartmann, "Photon Echoes" *Physical Review* **141**, 391–406, (1966).
(Cited on pages 58 and 59)

[102] G. J. Pryde, *Ultrahigh resolution spectroscopic studies of optical dephasing in solids*, Ph.D. thesis, Australian National University, (1999).
(Cited on pages 59 and 98)

[103] H. Lin, T. Wang and T. W. Mossberg, "Demonstration of 8-Gbit/in^2 areal storage density based on swept-carrier frequency-selective optical memory." *Optics Letters* **20**(15), 1658–60, (1995).
(Cited on page 60)

[104] J. Ruggiero, J.-L. Le Gouët, C. Simon and T. Chanelière, "Why the two-pulse photon echo is not a good quantum memory protocol" *Physical Review A* **79**, 053851, (2009).
(Cited on page 61)

[105] W. Tittel, M. Afzelius, T. Chaneliére, R. Cone, S. Kröll, S. Moiseev and M. Sellars, "Photon-echo quantum memory in solid state systems" *Laser & Photonics Reviews* **4**(2), 244–267, (2009).
(Cited on page 61)

[106] V. Damon, M. Bonarota, A. Louchet-Chauvet, T. Chanelière and J.-L. Le Gouët, "Revival of silenced echo and quantum memory for light" *New Journal of Physics* **13**, 093031, (2011).
(Cited on pages 32 and 61)

[107] G. Agarwal, *Quantum Statistical Theories of Spontaneous Emission and their Relation to Other Approaches*, Volume 70 of *Springer Tracts in Modern Physics*, Springer-Verlag, Berlin, (1974).
(Cited on pages 65, 70 and 189)

[108] E. T. Jaynes and W. Cummings, "Comparison of quantum and semiclassical radiation theories with application to the beam maser" *Proc. IEEE* **51**(89), (1963).
(Cited on page 66)

[109] H. Wiseman, *Complementarity in Spontaneous Emission: Quantum Jumps, Staggers and Slides*, in *Directions in Quantum Optics*, Edited by H. J. Carmichael, R. J. Glauber, M. O. Scully, Lecture Notes in Physics, **561**, 347, Springer (2001).
(Cited on page 68)

[110] R. Dicke, "Coherence in spontaneous radiation processes" *Physical Review* **93**(1), 99–110, (1954).
(Cited on pages 69, 75, 125 and 193)

[111] R. Glauber, "Coherent and Incoherent States of the Radiation Field" *Physical Review* **131**(6), 2766–2788, (1963).
(Cited on pages 71, 202 and 203)

[112] F. Arecchi, E. Courtens, R. Gilmore and H. Thomas, "Atomic coherent states in quantum optics" *Physical Review A* **6**(6), 2211, (1972).
(Cited on page 71)

[113] R. Friedberg and J. Manassah, "Dicke states and Bloch states" *Laser Physics Letters* **4**(12), 900–911, (2007).
(Cited on page 71)

[114] R. Glauber, "The Quantum Theory of Optical Coherence" *Physical Review* **130**(6), 2529–2539, (1963).
(Cited on pages 74, 199 and 202)

[115] J. H. Eberly, "Emission of one photon in an electric dipole transition of one among N atoms" *Journal of Physics B: Atomic, Molecular and Optical Physics* **39**(15),

S599–S604, (2006).
(Cited on pages 78 and 79)

[116] M. Scully, E. Fry, C. Ooi and K. Wódkiewicz, "Directed Spontaneous Emission from an Extended Ensemble of N Atoms: Timing Is Everything" *Physical Review Letters* **96**, 010501, (2006).
(Cited on pages 78 and 79)

[117] N. Treps and C. Fabre, "Criteria of quantum correlation in the measurement of continuous variables in optics" *Laser Physics* **15**(1), 187–194, (2005).
(Cited on page 86)

[118] T. S. Iskhakov, E. D. Lopaeva, a. N. Penin, G. O. Rytikov and M. V. Chekhova, "Two methods for detecting nonclassical correlations in parametric scattering of light" *JETP Letters* **88**(10), 660–664, (2009).
(Cited on pages 86 and 87)

[119] A. Heidmann, R. Horowicz, S. Reynaud, E. Giacobino, C. Fabre and G. Camy, "Observation of Quantum Noise Reduction on Twin Laser Beams" *Physical Review Letters* **59**(22), 2555–2557, (1987).
(Cited on page 87)

[120] M. D. Eisaman, *Generation, Storage, and Retrieval of Nonclassical States of Light using Atomic Ensembles*, Ph.D. thesis, Harvard University, (2006).
(Cited on pages 87, 88 and 92)

[121] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, Wiley, (1991).
(Cited on pages 88, 139, 202, 204 and 209)

[122] R. Loudon, "Non-classical effects in the statistical properties of light" *Reports on Progress in Physics* **43**(7), 913, (1980).
(Cited on page 89)

[123] T. Mossberg, A. Flusberg, R. Kachru and S. Hartmann, "Tri-Level Echoes" *Physical Review Letters* **39**(24), 1523–1526, (1977).
(Cited on page 99)

[124] G. J. Pryde, M. J. Sellars and N. B. Manson, "Solid State Coherent Transient Measurements Using Hard Optical Pulses" *Physical Review Letters* **84**, 1152–1155, (2000).
(Cited on pages 103 and 109)

[125] M. G. Benedict, A. M. Ermolaev, V. A. Malyshev, I. V. Sokolov and E. D. Trifonov, *Super-radiance Multiatomic Coherent Emission*, Taylor & Francis Group, (1996).
(Cited on page 125)

[126] L. Viola and S. Lloyd, "Dynamical suppression of decoherence in two-state quantum systems" *Physical Review A* **58**(4), 2733–2744, (1998).
(Cited on page 167)

[127] M. A. Pravia, N. Boulant, J. Emerson, A. Farid, E. M. Fortunato, T. F. Havel, R. Martinez and D. G. Cory, "Robust control of quantum information" *The Journal of Chemical Physics* **119**(19), 9993, (2003).
(Cited on page 168)

[128] E. Fraval, M. J. Sellars and J. J. Longdell, "Dynamic Decoherence Control of a Solid-State Nuclear-Quadrupole Qubit" *Physical Review Letters* **95**, 030506 (2005).
(Cited on page 170)

[129] J. J. Longdell, A. L. Alexander and M. J. Sellars, "Characterization of the hyperfine interaction in europium-doped yttrium orthosilicate and europium chloride hexahydrate" *Physical Review B* **74**, 195101, (2006).
(Cited on page 171)

[130] D. Meschede, *Optics, light and lasers: the practical approach to modern aspects of photonics and laser physics*, Wiley, (2004).
(Cited on pages 189 and 193)

[131] V. Weisskopf and E. Wigner, "Berechnung der natürlichen Linienbreite auf Grund der Diracschen Lichttheorie" *Zeitschrift für Physik A Hadrons and Nuclei* **63**(1), 54–73, (1930).
(Cited on page 189)

[132] M. Lukin, "Modern Atomic and Optical Physics II (Lecture Notes from Fall 2005)" (2005).
(Cited on page 192)

[133] R. Jodoin and L. Mandel, "Superradiance in an inhomogeneously broadened atomic system" *Physical Review A* **9**, 873–884, (1974).
(Cited on page 194)

[134] A. Walther, A. Amari, S. Kröll and A. Kalachev, "Experimental superradiance and slow-light effects for quantum memories" *Physical Review A* **80**, 012317, (2009).
(Cited on page 194)

[135] R. Hanbury Brown and R. Q. Twiss, "Correlation between photons in two coherent beams of light" *Nature* **4497**, 27–29, (1956).
(Cited on page 199)

[136] R. Hanbury Brown and R. Q. Twiss, "A test of a new type of stellar interferometer on Sirius" *Nature* **178**, 1046, (1956).
(Cited on page 199)

[137] E. Schrödinger, "Der stetige Übergang von der Mikro- zur Makromechanik" *Die Naturwissenschaften* **14**(28), 664–666, (1926).
(Cited on page 202)

[138] F. Steiner, "Schrödinger's discovery of coherent states" *Physica B+C* **151**(1-2), 323–326, (1988).
(Cited on page 202)

[139] R. Glauber, "Nobel Lecture: One hundred years of light quanta" *Reviews of Modern Physics* **78**(4), 1267–1278, (2006).
(Cited on pages 202 and 204)

[140] E. Sudarshan, "Equivalence of semiclassical and quantum mechanical descriptions of statistical light beams" *Physical Review Letters* **10**(7), 277, (1963).
(Cited on page 203)

[141] M. J. Sellars, *Ultra-High Resolution Laser Spectroscopy of Rare Earth Doped Solids*, Ph.D. thesis, Australian National University, (1995).
(Cited on page 213)

[142] S. Polyakov, A. Migdall and S. W. Nam, "Simple and Inexpensive FPGA-based Fast Time Resolving Acquisition Board". http://www.physics.nist.gov/fpga (accessed 2009).
(Cited on page 218)