

Chapter 3

A light-matter interface: quantifying coherence

One of the first experiments conducted in our group with a trapped intracavity atom was to demonstrate that it could generate single photons on demand by means of an adiabatic (dark-state) process [16]. After the ground-state cooling experiment, we turned our attention again to this process, with the idea of mapping cavity-generated photons back onto the atom. In this way, we could exploit an important feature of photon generation within a cavity: the fact that the output photon is created in a well-defined spatial mode. Moreover, the photon generation process is coherent and thus reversible: by running the same adiabatic process backwards, we can in principle achieve state transfer from light (photons) to matter (cesium hyperfine levels).

However, as we began to set up the experiment, we realized that using single photons would restrict our ability to characterize the phase coherence of the process we hoped to demonstrate. Coherence, meanwhile, is at the heart of the quantum information schemes for which this mapping could be employed. We decided instead to explore reversible state transfer by using a pulse of weak coherent light at the cavity input that had roughly the same temporal profile as a cavity-generated photon. From a pair of coherent pulses in sequence, we could then extract important phase information about the transfer process. The results of this experiment are reproduced in Section 3.1. The following sections are then devoted to a more thorough discussion of technical details, including the necessary pulse combinations and timing sequences,

the requirements of phase stability between pairs of laser pulses, and remote serial programming of electronic devices. I conclude by discussing the outlook for mapping single photons to and from the cavity in the future.

3.1 Reversible state transfer between light and a single trapped atom

The following section has been adapted from Ref. [31].

We demonstrate the reversible mapping of a coherent state of light with mean photon number $\bar{n} \simeq 1.1$ to and from the hyperfine states of an atom trapped within the mode of a high-finesse optical cavity. The coherence of the basic processes is verified by mapping the atomic state back onto a field state in a way that depends on the phase of the original coherent state. Our experiment represents an important step towards the realization of cavity QED-based quantum networks, wherein coherent transfer of quantum states enables the distribution of quantum information across the network.

An important goal in quantum information science is the realization of quantum networks for the distribution and processing of quantum information [1], including for quantum computation, communication, and metrology [44, 70, 71, 72]. In the initial proposal for the implementation of quantum networks [2], atomic internal states with long coherence times serve as “stationary” qubits, stored and locally manipulated at the nodes of the network. Quantum channels between different nodes are provided by optical fibers, which transport photons (“flying” qubits) over long distances [73]. A crucial requirement for such network protocols is the reversible mapping of quantum states between light and matter. Cavity quantum electrodynamics (QED) provides a promising avenue for achieving this capability by using strong coupling for the interaction of single atoms and photons [49].

Within this setting, reversible emission and absorption of one photon can be achieved by way of a dark-state process involving an atom and the field of a high-




Figure 3.1: Illustration of the protocol of Ref. [2] for quantum state transfer and entanglement distribution from system A to system B . By expanding to a larger set of interconnected cavities, complex quantum networks can be realized.

finesse optical cavity. For classical fields, this “STIRAP” process was first considered twenty years ago [74, 75], before being adapted to quantum fields [76] and specifically to the coherent transfer of quantum states between remote locations [2], with many extensions since then [77]. The basic scheme, illustrated in Figure 3.1, involves a three level atom with ground states $|a\rangle$ and $|b\rangle$ and excited state $|e\rangle$. An optical cavity is coherently coupled to the atom on the $b \leftrightarrow e$ transition with rate g , and a classical field $\Omega(t)$ drives the atom on the $a \leftrightarrow e$ transition. If the Ω field is ramped adiabatically $off \rightarrow on$, then state $|a, n\rangle$ evolves into $|b, n+1\rangle$, and state $|b, n\rangle$ remains unchanged, where $|a, n\rangle$, $|b, n\rangle$ denotes a state in which the atom is in ground state a , b and there are n photons in the cavity. Ramping Ω $on \rightarrow off$ implements the reverse transformation.

This process can be used to generate single photons by preparing the atom in $|a\rangle$ and ramping Ω $off \rightarrow on$, thereby effecting the transfer $|a, 0\rangle \rightarrow |b, 1\rangle$ with the coherent emission of a single photon pulse from the cavity [2, 76, 78]. Essential aspects of this process have been confirmed in several experiments [16, 79, 80], including tailoring of the single-photon pulse shape [79].

A distinguishing aspect of this protocol is that it should be *reversible* [2], so that a photon emitted from one system A can be efficiently transferred to another system B . Furthermore, it should be possible to map coherent superpositions reversibly between the atom and the field:

$$(c_0|b\rangle + c_1|a\rangle) \otimes |0\rangle \leftrightarrow |b\rangle \otimes (c_0|0\rangle + c_1|1\rangle). \quad (3.1)$$

Over the past decade, single photons have been generated in diverse physical systems [81]; however, most such sources are not reversible in principle, and for those that are, no experiment has verified the reversibility of either the emission or the absorption process.

We report an important advance related to the interface of light and matter by explicitly demonstrating the reversible mapping of a coherent optical field to and from the hyperfine ground states of a single, trapped Cesium atom.¹ Specifically, we map an incident coherent state with $\bar{n} = 1.1$ photons into a coherent superposition of $F = 3$ and $F = 4$ ground states with transfer efficiency $\zeta = 0.057$.² We then map the stored atomic state back to a field state. The coherence of the overall process is confirmed by observations of interference between the final field state and a reference field that is phase coherent with the original coherent state, resulting in a fringe visibility $v_a = 0.46 \pm 0.03$ for the adiabatic absorption and emission processes. We thereby provide the first verification of the fundamental primitive upon which the protocol in Ref. [2] is based.

As shown schematically in Figure 3.2(a), our system consists of one Cs atom coupled to a high-finesse Fabry-Perot cavity. The cavity length is tuned so that a TEM_{00} mode is near resonance with the $6S_{1/2}, F = 4 \rightarrow 6P_{3/2}, F = 3'$ transition of Cs at 852.4 nm. The maximum atom-cavity coupling rate is $g_0/2\pi = 16$ MHz, while the cavity field and the atomic excited state decay at rates $(\kappa, \gamma)/2\pi = (3.8, 2.6)$ MHz $\ll g_0$. Thus, the system is in the strong coupling regime of cavity QED [49].

Atoms are dropped from a magneto-optical trap into the cavity and cooled into a far off-resonant trap (FORT) by a blue-detuned optical lattice (see [49, 30]). The FORT excites another TEM_{00} cavity mode at the “magic” wavelength 935.6 nm, creating nearly equal trapping potentials for all states in the $6S_{1/2}, 6P_{3/2}$ manifolds [13].

¹This mapping could also be demonstrated by absorbing a single-photon state, but we use a coherent state instead, because its phase information allows us to verify explicitly the reversibility of the absorption process.

²Since we use a coherent state rather than a coherent superposition of $n = 0, 1$ Fock states, equation (3.1) only approximately describes our system. For a $\bar{n} = 0.68$ coherent state (i.e., $\bar{n} = 1.1$ at the face of M_{in}), the fraction of the population in the $n = 0, 1$ subspace is $\simeq 0.85$.




Figure 3.2: (a) Schematic of the experiment. The probe $\lambda(t)$ resonantly drives the cavity through input mirror M_{in} ; the classical field $\Omega(t)$ excites the atom transverse to the cavity axis. Photons emitted from the output mirror M_{out} are directed to a pair of avalanche photodiodes. (b) Atomic level diagram. Double arrow g indicates the coherent atom-cavity coupling, and $\Omega(t)$ is the classical field. The cavity and Ω field are blue-detuned from atomic resonance by Δ .

An atomic level diagram is shown in Figure 3.2(b); the states used in the current scheme are ground $F = 3, 4$ and excited $F = 3'$ manifolds, corresponding to $|a\rangle, |b\rangle, |e\rangle$ in Figure 3.1. The cavity is tuned to frequency $\omega_C = \omega_{4-3'} + \Delta$, where $\omega_{4-3'}$ is the frequency of the $4 - 3'$ transition, and $\Delta/2\pi = 10$ MHz is the cavity-atom detuning. A linearly polarized probe beam³ drives the cavity at frequency ω_C with pumping strength $\lambda(t)$. An optical lattice drives the atom transverse to the cavity axis at frequency $\omega_A = \omega_{3-3'} + \Delta$ to provide a classical field with Rabi frequency $\Omega(t)$ [16]. The laser source for the optical lattice is phase-locked in Raman resonance with the probe laser, so their relative detuning $\delta = \omega_A - \omega_C$ is phase-stable and equal to the ground-state hyperfine splitting $\Delta_{HF} = \omega_{3-3'} - \omega_{4-3'} = (2\pi)(9.193 \text{ GHz})$.

Our experimental procedure is as follows: after loading an atom into the FORT, we subject it to 2,000 trials lasting a total of 360 ms, where each trial consists of

³Our cavity supports two nearly degenerate modes with orthogonal linear polarizations along the \hat{x} and \hat{y} axes, where \hat{z} denotes the cavity axis. The cavity probe drives either the \hat{x} or the \hat{y} mode; a polarizer at the cavity output only allows the \hat{x} component to reach the detectors.

a series of discrete measurements performed on the atom. These measurements are used to quantify the coherence of the absorption process, as well as for calibrations and background monitoring. After these trials, we check that the atom has survived in the trap by attempting to generate 10,000 single photons, which are detected by monitoring the cavity output with two single-photon-counting avalanche photodiodes. We keep only the data from atoms that have survived all the trials. For most of the data that we keep, only a single atom is present in the trap, but occasionally two or more atoms may be loaded. From measurements performed during the 2,000 trials, we determine that at least 80% of the data presented here involve a single atom.

For each trial, we prepare the atom in $F = 4$ and then drive the system with a series of light pulses, as shown in Figure 3.3. The classical field $\Omega(t)$ generates pulses $\Omega_{1,2}$, and the cavity probe $\lambda(t)$ generates pulses $\lambda_{1,2}$. For any given measurement within a trial, some of these pulses are *on* and the others are *off*. Pulse λ_1 is the coherent state that is to be mapped into the atom. The strength of this pulse is set so that there are $\bar{n} = 1.1$ mode-matched photons at the face of the input mirror M_{in} . Because of mirror losses [42], if no atom were present, this would give rise to a pulse inside the cavity with $\bar{n} = 0.68$ photons. The falling edge of pulse Ω_1 is used to perform the adiabatic absorption of λ_1 . The intensity of the lattice light is such that when Ω_1 is fully *on*, its Rabi frequency is $\sim 8\gamma$, a value found to maximize the adiabatic absorption probability. When λ_1 is absorbed, some of the atomic population is transferred from $F = 4$ to $F = 3$. With λ_2 *off*, Ω_2 allows us to determine the fraction of the population that has been transferred: if the atom is in $F = 4$, then Ω_2 does nothing, while if the atom is in $F = 3$, then the rising edge of Ω_2 transfers it back to $F = 4$ and generates a single photon. Finally, with both pulses Ω_2 and λ_2 *on*, we verify that λ_1 was absorbed coherently. The Ω_2 , λ_2 pulses act together to generate a field inside the cavity; if λ_1 was absorbed coherently, then the amplitude of this field will depend on the relative phase θ between λ_1 , λ_2 .

This dependence can be understood by considering a simple model in which Ω_2 and λ_2 act independently. With λ_2 *off* and Ω_2 *on*, the Ω_2 pulse transfers the atom from a superposition of $F = 3, 4$ into $F = 4$ by generating a field α in the cavity




Figure 3.3: Timing diagram: the upper curve shows the Ω_1 and Ω_2 pulses; the lower curve shows the λ_1 and λ_2 pulses. Each of these pulses can be turned *on/off* independently. Here Δt is the delay between the falling edge of Ω_1 and the rising edge of Ω_2 . By enabling various combinations of these pulses, and/or varying the relative phase θ between λ_1 and λ_2 , we perform different measurements on the atom. The phase θ is set by the phase difference of RF pulses driving an acousto-optic modulator in the probe beam.

whose phase depends on the phase of the atomic superposition. In turn, the phase of the original atomic superposition is set by the phase of λ_1 . With λ_2 *on* and Ω_2 *off*, the λ_2 pulse generates a field β inside the cavity whose phase is set by λ_2 . If Ω_2 and λ_2 acted independently, then when both Ω_2 and λ_2 were *on*, the fields α and β would combine to give a total field $\alpha + \beta$, whose amplitude depends on the phase difference θ between λ_1 and λ_2 . Because Ω_2 and λ_2 do not act independently, this model is only approximately correct. Nevertheless, the phase of the final field still depends on θ for the coherent processes associated with $\lambda_{1,2}$, $\Omega_{1,2}$.

We first consider a series of measurements which demonstrate that the λ_1 pulse transfers more population from $F = 4$ to $F = 3$ in the presence of the Ω_1 pulse than in its absence. We start with the atom in $F = 4$ and apply the λ_1 pulse, either with the Ω_1 pulse (adiabatic absorption, which consists of both coherent and incoherent components) or without it (only incoherent absorption $4 \rightarrow 3'$, with spontaneous decay to $F = 3$). In either case, λ_1 transfers some population from $F = 4$ to $F = 3$.

To quantify the population transfer, we apply Ω_2 and measure the probability that a single photon is detected within $1 \mu\text{s}$ of the rising edge of Ω_2 .⁴ We thereby infer the fraction of the atomic population that was in $F = 3$.⁵ For adiabatic absorption (Ω_1 *on*), we find that the probability p_a for the atom to be transferred from $F = 4$ to $F = 3$ by λ_1 is $p_a = 0.063 \pm 0.002$, whereas for incoherent absorption (Ω_1 *off*), the probability is $p_i = 0.046 \pm 0.001$. The ratio of the adiabatic to the incoherent absorption probability is $r = p_a/p_i = 1.38 \pm 0.04$.

As shown in Figure 3.4, we vary the arrival time t_1 of the λ_1 pulse and study the effect on the adiabatic-to-incoherent ratio r .⁶ This ratio is maximized when λ_1 is well aligned with the falling edge of Ω_1 at $t = 0$. If λ_1 arrives too early ($t_1 \ll 0$), then any population that it transfers from $F = 4$ to $F = 3$ is pumped back to $F = 4$ by Ω_1 . If λ_1 arrives too late ($t_1 \gg 0$), then Ω_1 is already *off*, resulting in incoherent transfer with $r = 1$.

Figure 3.4 also shows the results of a computer simulation of the absorption process. The simulation predicts values for p_a and p_i and therefore the ratio $r = p_a/p_i$. The correspondence between our simulation and the actual measurements of r vs. t_1 in Figure 3.4 is qualitatively reasonable (the only free parameter in the simulation is the atom-cavity coupling g , which we set to $g/g_0 = 0.44$). The simulation can also be used to partition p_a into a coherent component p_a^c and an incoherent component p_a^i . We define the coherent component of r by $r^c = p_a^c/p_i$, the incoherent component of r by $r^i = p_a^i/p_i$, and plot r^c, r^i vs. t_1 in Figure 3.4. The simulation indicates that the value of t_1 for which the adiabatic absorption process is maximally coherent is roughly the value of t_1 that maximizes the adiabatic transfer probability, and suggests that for this value of t_1 the adiabatic absorption process has appreciable coherence, with $r^c/r^i \simeq 1$.

⁴For these measurements and those of Figure 3.5, $\Delta t = 290$ ns. The probe is polarized along \hat{y} ; for the case when λ_2 is *on*, this ensures that the emerging signal is not dominated by the component of λ_2 that is transmitted by the cavity.

⁵This involves subtracting a background probability of 0.0025, which we determine by pumping the atom to $F = 4$ and applying Ω_2 , and dividing by the single photon generation efficiency of 0.036, which we determine by pumping the atom into $F = 3$ and applying Ω_2 .

⁶For these measurements, $\Delta t = 2 \mu\text{s}$, and the probe is polarized along \hat{x} .




Figure 3.4: Ratio r of adiabatic transfer probability to incoherent transfer probability vs. arrival time t_1 for the incident coherent pulse λ_1 . Red data points (\circ): r vs. t_1 (experiment). Red solid curve: r vs. t_1 (computer simulation). Black dotted curve: coherent component r^c vs. t_1 (simulation). Blue dashed curve: incoherent component r^i vs. t_1 (simulation)

In Figure 3.5, we present measurements that demonstrate that the adiabatic absorption process is indeed coherent. As before, we prepare the atom in $F = 4$ and apply λ_1 , either with or without Ω_1 , followed by Ω_2 . But now we add the λ_2 pulse, which overlaps with the rising edge of Ω_2 . If the λ_1 pulse is absorbed coherently, then the amplitude of the field generated by the combined action of Ω_2 and λ_2 will depend on the relative phase θ of λ_1 and λ_2 . By recording the cavity output from M_{out} as a function of θ and observing this dependence, we can verify that the λ_1 pulse was absorbed coherently. To accomplish this, we repeat the above sequence for different values of θ , where for each relative phase, we measure the mean number of photons $n(\theta)$ emitted from the cavity within a fixed detection window. We take data both with Ω_1 *on* and *off*, so as to obtain results $n_a(\theta)$ and $n_i(\theta)$ both for adiabatic and incoherent absorption. Figure 3.5 plots $R_a(\theta) = n_a(\theta)/n_a(\theta_0)$ and $R_i(\theta) = n_i(\theta)/n_i(\theta_0)$, where θ_0 is a fixed phase. Note that these ratios, rather than the photon numbers themselves, are employed in order to cancel small, slow drifts in the intensity of the light beams. Significantly, we observe an appreciable phase-dependence with visibility




Figure 3.5: Ratios $R_a(\theta)$, $R_i(\theta)$ for photon generation as a function of the relative phase θ between the $\lambda_{1,2}$ fields. Red data points for adiabatic state transfer with Ω_1 *on*. Blue points for the incoherent process with Ω_1 *off*. The full curve is a fit to obtain the fringe visibility $v_a \simeq 0.46 \pm 0.03$. On average, each point represents about 130 atoms. The error bars represent statistical fluctuations from atom to atom.

$v_a = 0.46 \pm 0.03$ for the adiabatic absorption curve $R_a(\theta)$, while no such variation is recorded for the incoherent absorption curve $R_i(\theta)$.

The fringe visibility is limited by the intrinsic incoherent component of the absorption process, as well as by the mismatch in amplitudes and pulse shapes for the α , β fields. For the results shown in Figure 3.5, a 200 ns detection window is used around the peak of the emission process. If we increase the detection window to 1 μ s, thus degrading the pulse shape overlap, the visibility drops to $v_a = 0.18 \pm 0.01$.

In conclusion, we have demonstrated the reversible transfer of a coherent pulse of light to and from the internal state of a single trapped atom, which represents a significant step towards the realization of quantum networks based upon interactions in cavity QED. Explicitly, we have presented a detailed investigation of the adiabatic absorption of an incident coherent state with $\bar{n} = 1.1$ photons. A fraction $p_a = 0.063$ of the atomic population has been transferred from $F = 4$ to $F = 3$, with the efficiency of the transfer being $\zeta \equiv p_a/\bar{n} = 0.057$. Here ζ provides an estimate of the efficiency that could be obtained if we adiabatically absorbed a single photon state instead of

a coherent state, and should be compared to the much lower efficiencies possible in free space.

The factors that limit the transfer efficiency include the passive mirror losses [42], the fact that our cavity mirrors M_{in}, M_{out} have equal transmission coefficients $T_{in} = T_{out}$ (as opposed to $T_{in} \gg T_{out}$ for a single-sided cavity), and the coupling of the atom to both polarization modes of the cavity. Even in the ideal case without scatter and absorption losses in the mirrors, for a three-level atom coupled to a two-sided cavity ($T_{in} = T_{out}$) with two modes, the maximum possible adiabatic transfer probability would be $\zeta = 0.25$. By implementing a single-sided cavity with losses as achieved in Ref. [8], we estimate that ζ could be improved to $\zeta \sim 0.9$ for coupling schemes with a single polarization.

3.2 Pulse combinations

The measurements described above consisted of trapping an atom and preparing it in one of two hyperfine manifolds, followed by combinations of the four pulses Ω_1 , Ω_2 , λ_1 , and λ_2 shown in Figure 3.3. Specifically, we employed eight configurations in which Ω_2 was always on, the other three pulses were either on or off, and the phase between λ_1 and λ_2 was either fixed or varied (or not applicable, in the case of only one λ pulse). These possibilities are enumerated in Table 3.1. Note that during a single 360 ms trapping interval, an atom would be subject to all eight pulse combinations 2000 times, with appropriate preparation of the atomic state preceding each combination. This section focuses on the information that we are able to extract from each set of pulses.

Cases 1 and 2 in Table 3.1 correspond to times $t_1 \gg 0$ and $t_1 = 0$ in Figure 3.4. In case 1, the atom is prepared in $F = 4$; since the cavity is tuned to $F = 4 \rightarrow F' = 3$, we expect the vacuum Rabi splitting to shift the incoming pulse λ_1 out of resonance with the system. In the absence of Ω_1 , λ_1 will only transfer population to $F = 3$ through incoherent, off-resonant excitation of the atom. The probability of this population transfer is measured by attempting to generate a single photon using Ω_2 , which will

#	F_{start}	Ω_1	λ_1	λ_2	θ	measured probability
1	4	off	on	off	n/a	incoherent absorption
2	4	on	on	off	n/a	coherent and incoherent absorption
3	4	on	on	on	θ_0	adiabatic fringe, fixed phase
4	4	on	on	on	θ	adiabatic fringe, variable phase
5	3	off	off	off	n/a	single photon generation
6	4	off	off	off	n/a	background
7	4	off	on	on	θ_0	incoherent fringe, fixed phase
8	4	off	on	on	θ	incoherent fringe, variable phase

Table 3.1: The series of eight pulse combinations applied to each atom in the reversible state transfer experiment. F_{start} is the hyperfine manifold in which the atom is prepared before the pulses are applied.

succeed only if the atom is in $F = 3$. In case 2, the inclusion of a classical pulse Ω_1 means that now we expect the coherent Raman process to transfer population from $F = 4$ to $F = 3$, though of course incoherent excitation will still be possible. The relative probability to generate photons in these two cases reflects the extent to which our mapping process is coherent. While our system is in the strong coupling regime, we are limited by the fact that g is not larger, which means that some off-resonant light from λ_1 can enter the cavity. The maximum atom-cavity coupling for our system on the $F = 4 \rightarrow F' = 5$ transition is $g_{max} = 33.9$ MHz, but the single-photon generation scheme constrains us to the $F = 4 \rightarrow F' = 3$ transition ($g_0 = 16$ MHz), and the variation of g over the range of FORT wells and atomic motion result in an inferred coupling $g_0 = 7$ MHz, only a factor of two greater than the cavity linewidth.

Cases 3 and 4 are the basis for the data presented in Figure 3.5. We map the coherent pulse λ_1 onto the atom by means of Ω_1 ; then, while attempting to generate a single photon with Ω_2 , we introduce the pulse λ_2 . λ_2 interferes either constructively or destructively with the photon generation process as a function of θ , which in case 4 is varied over 2π . In case 3, θ is held fixed at an arbitrary value θ_0 , so the photon generation probability should in principle be constant. We compensate for small drifts in the probe power over the course of the experiment by expressing the phase variation as the ratio of the case 4 to case 3 data for each atom.

Case 5, in which we simply use Ω_2 to generate single photons on demand, provides us with a calibration of the photon generation efficiency. In case 6, we again attempt to generate single photons, but with the atom in the wrong hyperfine manifold, so that we expect nothing to happen. We then subtract this measured background value from the data obtained in all other cases.

Finally, cases 7 and 8 provide a baseline for comparison with cases 3 and 4. Because Ω_1 is omitted, we expect that the process that transfers population to $F = 3$ is strictly incoherent; that is, it should have no phase dependence. We confirm this by varying the relative phase of λ_1 and λ_2 and observing no change in photon generation probability. These data are also plotted for comparison in Figure 3.5; as in cases 3 and 4, we divide the probability with variable phase by the probability with fixed phase for each atom to remove the effects of probe drift. This reflects a more general strategy of cycling repeatedly through the full range of experimental parameters rather than performing various measurements sequentially in the course of a data run. As a result, we are less sensitive to long-term drift and can obtain real-time information about the progress of the experiment. Another example of this strategy, increasing the phase θ in 20° steps between each atom loading event, is discussed in Section 3.5.

3.3 Timing

Reversible state transfer was the first lab 11 experiment that required more complicated timing than the ADwin alone could provide. The ADwin can produce pulses as short as 100 ns; however, the length of any pulse shorter than a few μs should be confirmed on an oscilloscope, as it may be up to 1 μs shorter or longer than specified by the control program. In the present case, we had two requirements: (a) we needed to generate a pair of adiabatic pulses $\Omega_{1,2}$ from the side of the cavity that would turn on and off over hundreds of ns, and (b) we needed the timing of these pulses to be stable with respect to $\lambda_{1,2}$, the pair of coherent state pulses along the cavity axis. (We did not take special care to generate $\lambda_{1,2}$ adiabatically, as we found that the temporal pulse shape that resulted from switching the probe AOM with TTL logic

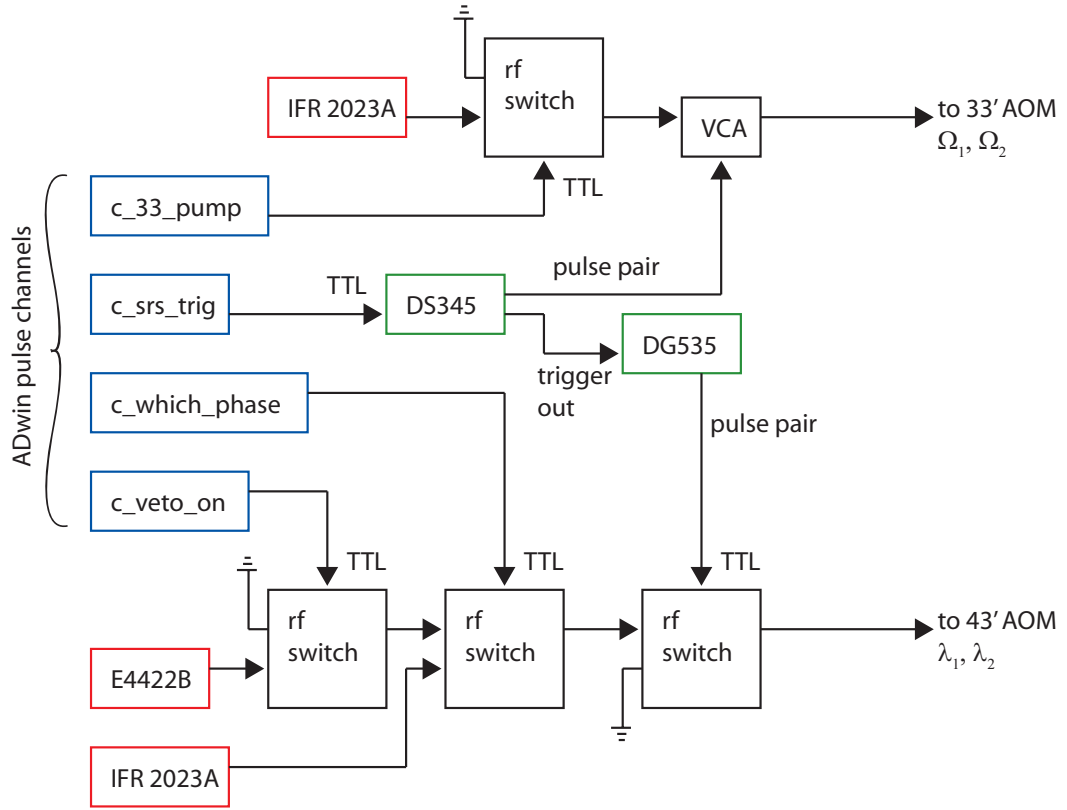


Figure 3.6: Instrumentation diagram for the generation of $\Omega_{1,2}$ and $\lambda_{1,2}$ pulses

was a reasonable approximation to the shape of photons generated in Ref. [16].)

As shown in Figure 3.6, timing was controlled by four digital logic output channels on the ADwin. Channel `c_srs_trig` triggered a single sweep of a Stanford Research Systems (SRS) DS345 function generator. The DS345 can be programmed to generate arbitrary waveforms; in this case, we set the output to shape the two Ω pulses shown in Figure 3.3. The pulse pair then provided the control input for a voltage-controlled attenuator (VCA). A second channel, `c_33_pump`, controlled the RF input to the VCA and thus regulated whether one or both of the pulses would drive the $F = 3 \rightarrow F = 3'$ AOM. Meanwhile, the $\lambda_{1,2}$ pulse pair was generated by a pre-programmed SRS DG535 pulse generator using the logical AND of its two pulse outputs. A single sweep of the DG535 was triggered by the DS345 rather than directly by the ADwin in order to avoid ~ 25 ns jitter between the pulse pairs due to the DS345 clock

rate. Two other ADwin channels controlled RF switches in the $\lambda_{1,2}$ path: the first, `c_which_phase`, determined which of two signal generators would provide the RF source; both signal generators were set to the same frequency, but one (IFR 2023A) had a fixed phase, while the phase of the other (Agilent E4422A) was variable. A second switch, `c_veto_on`, performed the same function as `c_33_pump`, determining whether none, one, or both of the pulses would be able to pass to the $F = 3 \rightarrow F = 4'$ AOM.

3.4 Phase stability

Our first attempts to generate a fringe using both $\Omega_{1,2}$ and $\lambda_{1,2}$ were not successful, though the data suggested some underlying phase-dependent periodicity. We realized that the Raman process underlying photon absorption and generation relied on the relative phase of $\Omega(t)$, the $F = 3 \rightarrow F' = 3$ lattice, and $\lambda(t)$, the source of $F = 3 \rightarrow F' = 4$ pulses along the cavity axis. While $\lambda(t)$ was derived from the master (probe) laser, $\Omega(t)$ was generated by the repump laser, which at the time had a separate lock to the cesium saturated-absorption spectrum. Luckily, the electronic phase lock that we had formerly used to synchronize the FORT and Raman lasers had been decommissioned during the ground state cooling experiment in favor of a quieter injection lock (Section 2.2.1). We resurrected it in order to set up a phase lock between the master laser and the repump: auxiliary beams from both lasers were combined on a fiber beamsplitter whose output was focused onto a 25 GHz photodetector (New Focus 1417). In the same configuration used previously for the Raman lock, the optical beat note from these two lasers was mixed down at Δ_{HF} and provided feedback signals for the repump laser's diode current and piezo voltage.

The fringe measurement also depended on a stable, programmable phase relationship θ between λ_1 and λ_2 . The two signal generators for λ_1 and λ_2 (Figure 3.6) were set to provide the same output frequency and were phase-referenced to one another using a 10 MHz input/output. In order to confirm relative phase stability, we monitored the outputs of the two sources (set to a fixed relative phase) over the course

of two days on a fast oscilloscope and found the long-term phase drift to be between 0.02 and 0.05 radians/hour. The data shown in Figure 3.5 were acquired over the course of about three hours, so we would expect drifts of $< 0.15 \text{ rad} \approx 2 \text{ degrees}$ during this time frame.

A second verification of the phase stability was provided by an optical interferometer formed by the shifted (first-order) and unshifted (zeroth-order) beams at the output of the $F = 4 \rightarrow F' = 3$ AOM, driven by one of the two RF sources. We generally use an iris to block all but the first-order beam, but with the iris opened, the two beams could beat against one another when they were recombined after being double-passed through the AOM. The beat note, at twice the RF input frequency, was then observed to be phase-stable on an oscilloscope with respect to the second RF source.

3.5 Remote programming capabilities

One important technical development for this experiment was the ability to program electronic equipment remotely via serial commands sent from the ADwin. During the ground-state cooling experiment, we had adjusted the Raman frequency detuning by hand on an IFR (Aeroflex) 2023A signal generator. That process grew increasingly time consuming, however, as our experiments became more complicated.

We now include ASCII strings as text in the code that controls our experiment. For example, the instructions

```
0.0 'text = ":PHASe "'
0.0 inline 'serial_partial_out (9)'

0.0 inline 'lngtostr (10*(i-1), text)'
0.0 inline 'serial_partial_out (9)'

0.0 inline 'text = " DEg"'
0.0 inline 'serial_out (9)'
```


tell the Agilent E4422B signal generator at time 0.0 to set the phase of its output frequency to $10 * (i - 1)$ degrees, where i is an integer between 1 and 18. (As the RF signal drives a double-passed AOM, an RF phase shift of 180° is equivalent to an optical phase shift of 360° .) The subroutines `serial_partial_out` and `serial_out` send the ASCII characters as TTL pulses on a specified ADwin channel (in this case, channel 9), followed by either a carriage return or a line feed to signal the end of the command. A home-built circuit based on a Texas Instruments MAX232N chip then converts the TTL pulses to signals on an RS-232 serial cable. The serial cable can be used either to drive devices directly (for example, the Aeroflex and Agilent signal generators) or to drive GPIB-controlled devices such as the DG535. In the latter case, we use an RS232-to-GPIB converter manufactured by National Instruments.

In the process of data acquisition, we run a loop which increments i on an atom-by-atom basis. In addition to writing time-stamped photon counts to two channels of the P7888 card, we write the value of i to a third channel. The data analysis programs that David Boozer has written then separate the atoms into bins indexed by the value of i .

Our first remote programming applications in the lab were automated frequency sweeps (in order to record Raman spectra) and phase shifts (to observe the coherence fringe in the current experiment). More recently, we have generated pulse trains of variable lengths and delays in order to measure, for example, Rabi flops and Ramsey fringes. We have also extended our frequency sweeps to the microwave domain, where they enable us to search for narrow resonance features (Chapter 5).

3.6 Mapping single photons

We return now to consider the original plan to map cavity-generated photons back onto a trapped atom. Figure 3.7 outlines a scheme that would allow us to realize this in the laboratory. A photon exiting the cavity passes (with probability 1/2) through a nonpolarizing beamsplitter cube (NPBS), a Faraday rotator, and a PBS into an optical delay line, which would provide a necessary buffer for the experiment timing.




Figure 3.7: Possible output path configuration in order to generate single photons on demand, then map them back into the cavity

At the fiber output a few microseconds later, a half-waveplate rotates the photon so that it reflects off the PBS and retraces its earlier path. A second pass through the Faraday rotator corrects for the waveplate rotation, and the photon re-enters the cavity in its original polarization state.

Half of all photons at the cavity output reflect off the NPBS, where they can be detected by one or more SPCMs; this includes both photons generated within the cavity and returning photons that are not absorbed. The initial step to characterize photon reabsorption into the cavity would be to monitor returning reflected photons at this port. As the timing of the adiabatic classical field Ω_1 is adjusted so that it is synchronous with the returning photons, we would expect to see a reflection dip. One could then attempt to confirm photon reabsorption directly by generating a second photon in the cavity, which would only occur if the first photon had transferred the atom back to its initial hyperfine state.

It is important to acknowledge the inefficiency of this scheme, which is nevertheless the most promising one we have found. Table 3.2 enumerates losses at each stage of the proposed experiment after generation of an initial photon within the cavity mode, which we assume occurs on every attempt. We see that detection of a second photon

Propagation step	Efficiency
cavity escape (mirror losses)	0.6
symmetric cavity	0.5
NPBS	0.5
Faraday rotator	0.93
free space propagation (filters)	0.66
single-mode fiber coupling	0.7
propagation through 250 m single-mode fiber	0.75
polarization drift in fiber	0.98
Faraday rotator, second pass (includes polarization error)	0.89
NPBS	0.5
after reflection from cavity, NPBS	0.5
free space propagation (filters)	0.66
single-mode fiber coupling	0.7
SPCM quantum efficiency	0.5
probability to detect reflected photons	0.0024
cavity entrance (mirror losses)	0.6
symmetric cavity	0.5
cavity escape (mirror losses)	0.6
probability to detect a second photon	0.0004

Table 3.2: Path efficiencies in mapping a photon back into the cavity. We consider detection via both reflection from the cavity (in the absence of a synchronous classical mapping field) and generation of a second photon upon successful mapping.

would be roughly 100 times less efficient than single photon generation. It would thus be a challenge to optimize various parameters of the experiment (for example, timing of Ω_1) based on detection efficiency. In future cavity systems, of course, the possible changes mentioned in Section 3.1 — namely, the use of a single-sided cavity with reduced scattering and absorption losses — would improve the efficiency with every pass in or out of the cavity.

Other prospects include mapping cavity-generated photons onto the ensemble of cesium atoms in the lab 2 experiment, thus generating entanglement between the hyperfine states of our trapped atom and the atomic ensemble. In this case, however, we would prefer to have the atom’s hyperfine states entangled with photon polarization states rather than the Fock states ($|0\rangle, |1\rangle$), as polarization states are more robust to path losses. Schemes for cavity-QED generation of polarization-entangled photons

[4, 82, 83] require the ability to prepare the atom in a specific Zeeman sublevel, something that has been a challenge for us in the past. Chapter 4 describes our recent progress in Zeeman state preparation of single atoms, while Chapter 5 discusses our subsequent characterization of Raman transitions between Zeeman levels, including steps toward polarization entanglement.