FPGA-Controlled Laser Bank for the Engineering of Photonic Materials

Lindsay Bassman

June 2, 2014

Contents

List of Figures

List of Tables

1 Introduction

Photons are massless, uncharged particles, which at low density do not ordinarily interact with one another. As such, light beams in free space can intersect without any observable interaction. When light enters a medium, it can be attenuated or phase-shifted by amounts depending on the optical properties of the medium. In some media, these changes depend linearly on the intensity of the light, whereas in others the relationship is nonlinear. Media with nonlinear optical properties have the ability to mediate interactions between photons. For

this reason, much attention is now being focused on engineering materials that exhibit single photon nonlinearity, where the addition of a single photon changes the optical properties of the material. Single photon nonlinearity holds great promise for applications in quantum computing.

Quantum computing involves the processing of qubits, two-level quantum mechanical structures that hold one bit of information, as opposed to the classical bits ours computers use today. While bits are decidedly either 0 or 1 at any given time, qubits can exhibit superpositions of states and entanglement between one another. The special quantum characteristics of qubits would allow for enhanced execution speed of classically hard problems, like searching unordered lists and factorization, which current computers can only tackle with extremely time and energy expensive brute force methods. An extremely exciting prospect of building a quantum information processor is that it would have the ability to simulate quantum many-body systems, something much too complex for classical computers to model. Simulation of many-body systems would open a unique window into the quantum world, answering some questions while raising new ones.

Photons are ideal candidates for qubits as they are fast, cheap, easy to guide, and most importantly have weak free-space interactions, meaning the information the photon holds is robust to corruption through interaction with the environment. In order to be used as qubits, however, photons must be able to interact, or communicate, with one another. As mentioned above, media that possess strong single photon nonlinearity are able to mediate interactions between photons, and thus are integral to creating a quantum computer with photonic qubits. In our experiment, we aim to engineer just such a material, from the bottom up, out of Rydberg polaritons and study its properties and potential use in quantum information processing.

Building this synthetic material begins with a gas of Rubidium atoms that are cooled and trapped in a magneto-optical trap (MOT) in an ultra-high vacuum chamber. Next, a thin slice of the atoms is transported to an optical cavity via an atomic conveyor belt. Once in the cavity, multiple lasers will be used to create coupled excitations of light and a group of atoms, called Rydberg polaritons. Ultimately, we hope to observe strong, effective photon-photon interactions through these Rydberg polaritons. Exquisite control of multiple lasers will be necessary to carry out the entire process. Precise frequency output, a large range of frequencies, fast frequency jumps, and smooth, well-timed frequency ramps are among the requirements. In order to meet these needs I have developed a field programmable gate array (FPGA) to control up to four lasers simultaneously, providing time resolution in the couple hundred nanoseconds, a frequency range over 2 GHz, and frequency resolution in the picohertz.

2 Theory

2.1 Optical Nonlinearity

Optical nonlinearity is observed through changes in the optical properties of media as the intensity of incident light changes. It arises due to the dipole moment per volume, or polarization, of the material depending nonlinearly on the strength of the incident electromagnetic field. In linear optics, the polarization depends linearly on the applied electric field and can be written

$$P(t) = \chi E(t) \tag{1}$$

where χ is a complex value called the linear susceptibility of the material. The real part of χ corresponds to phase shifts of the field, while the imaginary part accounts for amplitude changes. Nonlinear optics can be modeled by expanding

$P(t)$ in a power series in the electric field strength as

$$P(t) = \chi_1 E(t) + \chi_2 E(t)^2 + \chi_3 E(t)^3 + \dots \tag{2}$$

For an atomic medium the dominant nonlinear term is χ_3, the third order nonlinearity, known as the optical Kerr nonlinearity. In order to observe photon-photon interactions, we require the system to exhibit single photon nonlinearity, a regime where adding just one more photon is enough to change the optical properties of the medium. For this to occur, the Kerr nonlinearity must be above a certain threshold. In conventional optical media, such as air or water, the Kerr nonlinearity is of order 10^{-25} and 10^{-22}, respectively. These values are ~ 20 orders of magnitude too small for single photon nonlinear optics (1). Using light that is resonant with an excitation of the medium can enhance the nonlinearity; however, it comes at the expense of large transmission attenuation. Electromagnetically induced transparency (EIT) allows us to circumvent this issue by rendering an optically thick medium transparent to a probe beam on resonance, thus avoiding absorption.

2.2 EIT

EIT is a coherent phenomenon that arises in a three level quantum system coupled by a weak probe field and strong control field. Through the creation of a dark state, EIT makes a resonant, opaque medium transparent via quantum interference. To conceptualize: consider a three-level ladder system with a ground state $|g\rangle$, an intermediate excited state $|e\rangle$, and a Rydberg state $|r\rangle$, see Figure[1].

In this scheme, $|g\rangle$ and $|e\rangle$ are separated by an energy $\hbar\omega_{eg}$; $|e\rangle$ and $|r\rangle$ by an energy $\hbar\omega_{er}$. The probe laser, at frequency ω_p, drives the transition $|g\rangle \rightarrow |e\rangle$ with detuning $\Delta_p = \omega_p - \omega_{ge}$ and Rabi frequency Ω_p, while the control laser, at

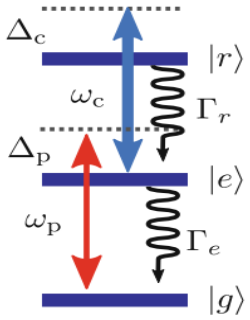

Figure 1: A three-level ladder system including a ground state, intermediate excited state, and Rydberg state.

frequency ω_c, drives the transition $|e\rangle \rightarrow |r\rangle$ with detuning $\Delta_c = \omega_c - \omega_{er}$ and Rabi frequency Ω_c. We can succinctly express the eignestates of the Hamiltonian of this system in terms of mixing angles θ and ϕ, defined by:

$$tan(\theta) = \frac{\Omega_p}{\Omega_c} \tag{3}$$

$$tan(2\phi) = \frac{\sqrt{\Omega_p^2 + \Omega_c^2}}{\Delta_p} \tag{4}$$

The eigenstates, derived in (2) are:

$$|+\rangle = sin(\theta)sin(\phi)\,|g\rangle + cos(\phi)\,|e\rangle + cos(\theta)sin(\phi)\,|r\rangle \tag{5}$$

$$|D\rangle = cos(\theta)\,|g\rangle - sin(\theta)\,|r\rangle \tag{6}$$

$$|-\rangle = sin(\theta)cos(\phi)\,|g\rangle - sin(\phi)\,|e\rangle + cos(\theta)cos(\phi)\,|r\rangle \tag{7}$$

In the EIT scheme, we require the probe beam to be much weaker than the control beam, in which case mixing angle $\theta \rightarrow 0$. On resonance ($\Delta_p = 0$), we then get the dressed eigenstates $|\pm\rangle = \frac{|r\rangle \pm |e\rangle}{\sqrt{2}}$ and $|D\rangle = |g\rangle$. The probe laser couples to the $|e\rangle$ components of the $|\pm\rangle$ states with equal and opposite amplitudes, resulting in destructive interference of this excitation path. Thus,

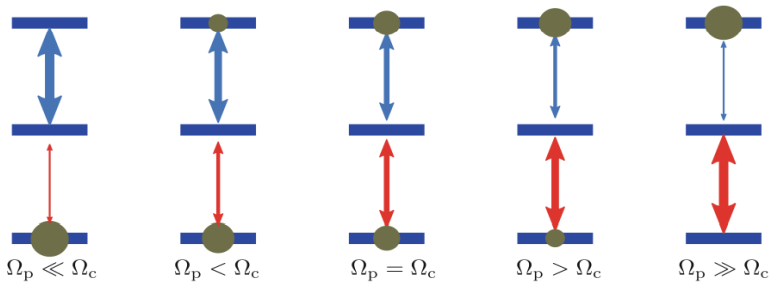

Figure 2: Dark state population distribution as a function of Ω_p and Ω_c. Initially all the atoms reside in the ground state when $\Omega_p << \Omega_c$. Adiabatically increasing Ω_p while decreasing Ω_c leads to a coherent population transfer (CPT) of the atoms into the Rydberg state, a process called STIRAP.

the probe beam is no longer absorbed, creating a window of transparency. The ground state is now the aforementioned dark state, since it does not couple to the probe beam. By adiabatically increasing Ω_p, while decreasing Ω_c we can coherently transfer atoms from the ground state to the Rydberg state, which will now be the dark state, depicted in Figure[2]. This was first demonstrated by Bergmann *et. al* (3), who dubbed this process stimulated Raman adiabatic passage (STIRAP).

Amazingly, EIT can actually be observed in an optical cavity without a control beam, which is replaced by having the mode of the cavity tuned to the $|e\rangle$ to $|r\rangle$ transition. When the probe beam is sent in, it is initially absorbed, but photons are preferentially emitted at the cavity's resonant frequency. These photons then perform the same function as the control beam in EIT. The probe beam thus induces its own transparency in a phenomenon called vacuum induced transparency, first theorized by Field in 1992 and experimental realized by Tanji-Suzuki *et. al* in 2011 (4). A resonant beam, together with EIT, can increase the Kerr nonlinearity by up to 16 orders of magnitude, though it still falls short of single photon nonlinearity.

2.3 Cooperativity

In order to exhibit single photon nonlinearity, we abandon processes mediated by the interaction between single quanta of photons and matter, and turn to processes involving the interaction between light-induced excitations. Such interactions are cooperative processes that require the collective, coherent response of a group of atoms to light. Cooperative phenomena will occur in systems where the interatomic interactions dominate the evolution of the ensemble. For an ensemble of atoms in a cavity, the figure of merit for cooperativity of atoms is derived from the coupling between the cavity and a single atom, parameterized by the single atom cooperativity η. η essentially measures the probability that an atom will scatter light into a cavity mode as opposed to into free space. For a cavity tuned to the atomic resonance, η is defined by

$$\eta = \frac{4g^2}{\kappa\lambda} \tag{8}$$

where g is the coupling strength of the atom to the cavity, λ is the natural linewidth of the atom, and κ is the cavity linewidth. When an ensemble of N atoms is placed in a cavity, cooperativity is enhanced by a factor of N. Thus, even if single atom cooperativity is low, with a large ensemble we can attain a large cooperativity, $N\eta$, for the system.

2.4 Dipole Blockade

We will be mainly interested in the cooperative behavior that arises from atoms that experience strong dipole-dipole interactions between one another. In this case, the properties of each atom will now depend on the presence of neighboring atoms, giving rise to cooperativity amongst the atoms. Rydberg atoms, with large orbital radii and millimeter spacing between Rydberg levels experience

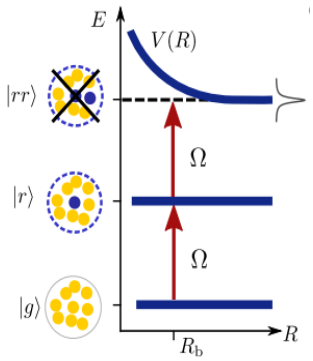

Figure 3: Schematic diagram depicting the dipole blockade phenomenon. For an ensemble of N atoms in the ground state $|g\rangle$ within the volume of a sphere of blockade radius R_b, one atom can be excited to a Rydberg state, transforming the ensemble to state $|r\rangle$. However, due to the position dependent energy level shift $V(R)$ of the doubly excited state $|rr\rangle$, a second Rydberg excitation is disallowed within the radius R_b.

strong dipole-dipole interactions which manifest as position dependent level shift $V(R)$. For the Rydberg atoms in the optical cavity we will have $V(R) = -\frac{C_6}{R^6}$. The cooperative phenomenon that arises from this position dependent level shift is known as dipole blockade, which prevents more than one Rydberg excitation from occuring within a sphere of blockade radius $R_b = (\frac{C_6}{\hbar\Omega})^{\frac{1}{6}}$, where Ω is the Rabi frequency of the resonant coupling between the ground and Rydberg states, shown schematically in Figure[3]. These blockade radii are typically on the order of a few microns. Since only one atom within the blockade radius can be excited to a Rydberg state, all atoms within this radius exhibit cooperative behavior, and can no longer be treated independently. This blockade allows deterministic creation of a singly excited collective state.

2.5 Laser Cooling and Trapping

The best way to attain a well-controlled and theoretically accessible system is to freeze out as many degrees of freedom as possible by cooling and trapping

atoms. Cold atoms are slow-moving, which limits the Doppler broadening when probing the atoms with lasers, providing more precise and efficient atom-photon interactions. Trapped atoms can be spatially located in a precise manner.

An effective way to decelerate and cool atoms depends on the scattering force on atoms moving in a laser beam, which arises from the momentum transfer between the atoms and the radiation field at or near resonance with an atomic transition. When an atom absorbs light, it stores the photon's energy, $\hbar\omega$, by going into an excited state and stores the photon's momentum, $\hbar k$, by recoiling from the light source with that momentum. The resultant change in velocity from recoil, $v_r = \frac{\hbar k}{M}$ will be on the order of a few cm/s. Although small compared to the thermal velocity of the atoms, multiple absorptions can ultimately result in a large total change in velocity. Since the atom can only absorb photons at certain frequencies, this method of deceleration is frequency dependent. Furthermore, the frequency of the laser light an atom sees will depend on it's velocity due to the Doppler effect, making this slowing force velocity dependent. It is because this force is velocity dependent that it can be used not only for deceleration, but for cooling that results in increased phase space density.

A straightforward implementation of this radiative deceleration is to have atoms travel in a diction opposite to the direction of a laser beam. In this case, atoms will absorb a photon, experiencing a small decelerating momentum kick from each photon. Between each absorption, the atom must return to its ground state through spontaneous decay, which involves emission of a photon in a random direction. The momentum transfers from such emissions will average to zero after many emission events. Therefore, the net momentum transfer to atoms after many interactions with photons will cause the atom to decelerate in the direction of the laser beam.

If we use two counter-propagating laser beams with the same frequency,

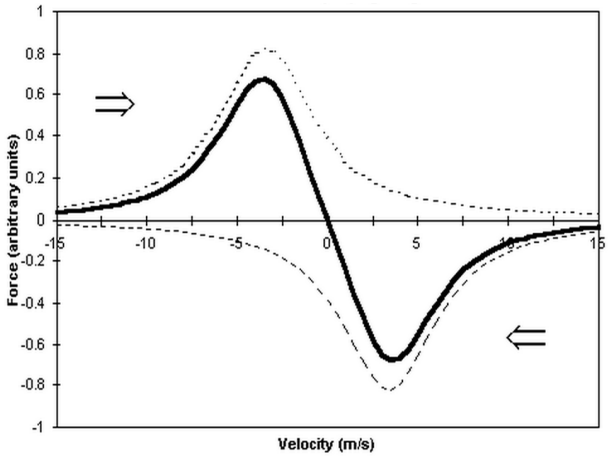

Figure 4: Velocity dependence of decelerating force on atoms in a one-dimensional optical molasses. The two dashed curves show the force of each laser beam, while the solid curve depicts the net force on the atom.

the atoms will experience a net force proportional to their velocity with a sign dependent of the laser frequency the atoms sees. If we detune the laser below the atomic transition, an atom moving opposite to the light will see a blue shifted frequency, which will be closer to resonance, and an atom moving with the light will see a red-shifted frequency farther from resonance. Therefore, atoms will interact more strongly with the beam that opposes their velocity and as a result be slowed, as if in an optical molasses (OM), the name aptly given to this technique. Figure[4] shows the force on the atom as a function of its velocity. OM can easily be extended to three dimensions with the use of three sets of mutually perpendicular counter-propagating lasers. Atoms will be slowed and confined in all three dimensions at the intersection of the six beams.

There is a limit to the atomic cooling, which is caused by heating of the atoms by the light beams. This heating derives from the quantum mechanical nature of the absorption and emission of photons, which dictates that the atom can only change its momentum in discrete steps of size $\hbar k$. This means the atom's kinetic energy can only change by a discrete amount, equal to at least the

recoil energy $E_r = \frac{\hbar^2 k^2}{2M} = \hbar\omega_r$. Therefore, the average frequency of absorption $\omega_{abs} = \omega_a + \omega_r$, while the average frequency of emission is $\omega_{em} = \omega_a - \omega_r$, where ω_a is the atomic transition resonant frequency. So we see that the light field will lose an average energy of $\hbar(\omega_{abs} - \omega_{em}) = 2\hbar\omega_r$ per scattering event. Due to conservation of energy, this energy must go somewhere, which in this case is into the atomic kinetic energy as the atoms recoil from each event. The atomic sample is heated by the random direction of the emission recoil. The steady state of atoms in OM arises when the competing rates of heating and cooling are equal, and will have a non-zero kinetic energy. The lower limit temperature found from the kinetic energy is called the Doppler temperature, T_D.

The velocity-dependent force from the lasers succeeds in cooling the atoms, confining them in velocity space. In order to spatially confine the atoms, we analogously require a spatially dependent potential for the atoms. We can accomplish this by employing both electric and magnetic fields to create a magneto-optical trap (MOT), first demonstrated in 1987 (5). The MOT leverages both an inhomogeneous magnetic field as well as radiative selection rules to achieve cooling and trapping through the radiative force and optical pumping. Counter-propagating laser beams of opposite circular polarization are used to set up the electric field. Anti-Helmholtz coils are used to set up a magnetic quadrupole field, as shown in Figure[5]. Due to the Zeeman effect, the presence of a magnetic field, will split the energies of the magnetic sublevels of an angular momentum state by adding an energy of $\mu_B B m_f$ to each of them (to first order). Thus, for $B > 0$, magnetic sublevels with $m_f > 0$ will have their energy shifted up, while sublevels with $m_f < 0$ will have their energy shifted down.

Trapping in the MOT is carried out by optical pumping of low velocity atoms in a linearly inhomogeneous magnetic field. To see how, we consider the simple scheme of pumping the atomic transition $J_g = 0 \rightarrow J_e = 1$ in one dimension,

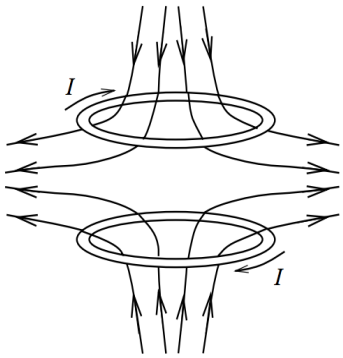

Figure 5: Two axially aligned coils with currents circulating in opposite directions and the resultant magnetic fields lines. The equal and opposite currents lead to a magnetic field with a magnitude of zero at the center of the system.

illustrated in Figure[6]. Here the $J_e = 1$ state has three magnetic sublevels: $M_e = -1, 0, 1$. Two counter-propagating lasers with opposite circular polarization, each detuned below the zero field atomic resonance by δ, are incident on the atoms. At position z', shown in Figure[6], the $M_e = -1$ energy level is shifted down, closer to resonance, while the $M_e = +1$ energy level is shifted up, farther out of resonance. Now selection rules dictate that a σ^+ polarized field couples states with $\Delta M = +1$ and a σ^- polarized field couples states with $\Delta M = -1$. Thus, an atom at position z', where $M_e = -1$ is closer to resonance, will scatter more light from the σ^- beam. An atom at position $-z'$, where the energy changes of $M_e = \pm 1$ states are reversed, will scatter more light from the σ^+ beam. For a MOT arranged as in Figure[6], with the σ^+ beam incident from the left and the σ^- beam incident from the right, the atoms will always be driven towards the center of the trap at $z = 0$. The physics is analogous to the velocity damping in optical molasses, except now instead of confining atoms in velocity space we are confining them in position space. The power of the MOT is that since the lasers are detuned below the atomic resonance, just as in OM, we can simultaneously achieve compression and cooling of atoms.

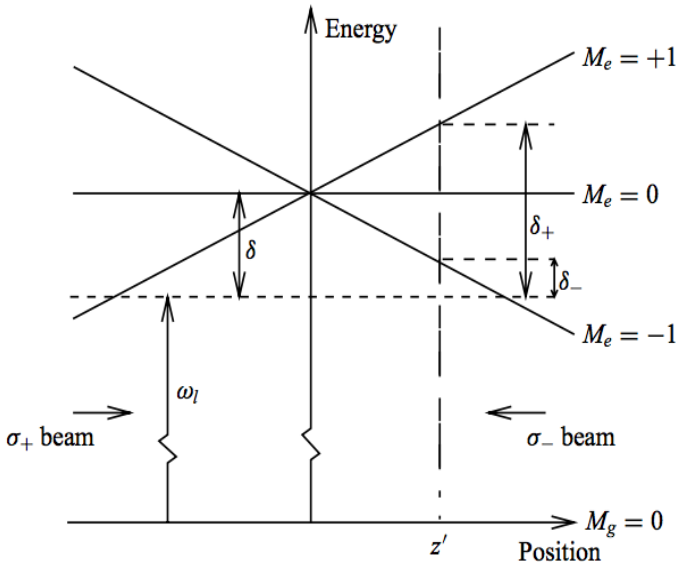

Figure 6: ...

Even better, since the laser beams are orthogonally, circularly polarized we can achieve cooling below the Doppler limit T_D, due to an effect called polarization gradient cooling introduced and explained by Dalibard and Cohen-Tannoudji (6). It is this polarization gradient cooling effect that allows us to achieve ultracold temperatures $\sim 10^{-6}K$. This MOT scheme can easily be extended to three dimensions with the use of six beams, and the scheme works for optical pumping of any transition where $J_g \rightarrow J_e = J_g + 1$. Together, the six beams and the coils, shown in Figure[7], will reduce the temperature of the Rubidium atoms and confine them into a small region of space, increasing density.

2.6 Atom Transport

After the atoms are cooled and compressed in the MOT, a thin slice of the atoms are transported via an atomic conveyor belt from the MOT into an optical cavity. The transport path is shown in Figure[8]. The conveyor belt is comprised of

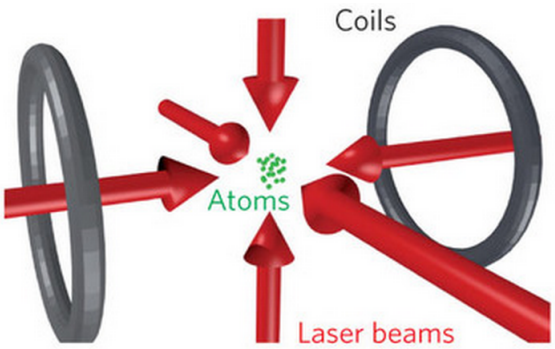

Figure 7: Three pairs of mutually perpendicular counterpropagating laser beams, together with two anti-Helmholtz coils slow and trap Rubidium atoms at their center.

a standing wave with a tight, moving focus. Fast focus translation is achieved with an objective mounted to a piezo linear actuator. Ray optics theory shows that a small translation of the objective results in a large translation of the focus. Once in the optical cavity, the coupling between an electromagnetic field and the trapped atoms is greatly enhanced due to repeated encounters of the beam with the atoms.

3 Precision Requirements

In order to accomplish the various steps required to build a cold, dense gas of Rydberg atoms and place it in a cavity, multiple fine-tuned lasers are necessary.

3.1 Specs for laser cooling and trapping

Our sample of Rubidium atoms will initially be in the ground state, $5^2S_{\frac{1}{2}}$, at about room temperature, which is $\sim 293K$, an energy of $\sim \frac{1}{40}eV$. The next levels the atoms can be excited to are $5^2P_{\frac{1}{2}}$, called the $D1$ line and $5^2P_{\frac{3}{2}}$, the $D2$ line, with their transition information shown in Figure[9]. From Figure[9]

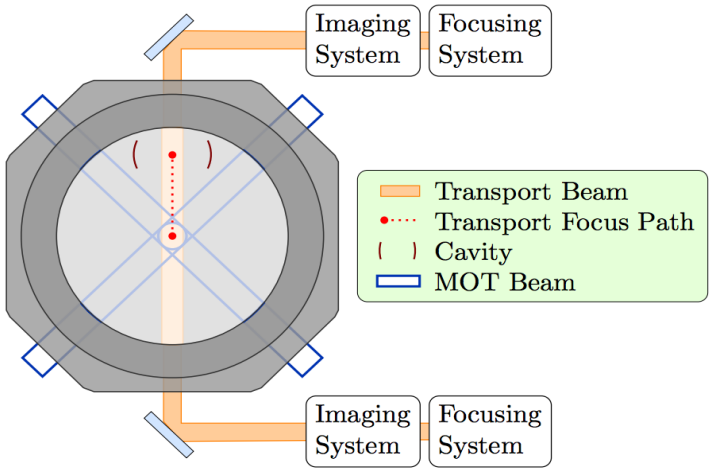

Figure 8: Atoms trapped in a MOT in the middle of a UHV chamber are transported via a standing wave with a moving focus into an optical cavity.

we see that it requires $\sim 1.5eV$ of energy to transition to either of the $5P$ states, so at room temperature the atoms are overwhelmingly likely to be in the $5S$ ground state. In the MOT, we will need to optically pump the atoms along the D2 line, $5^2S_{\frac{1}{2}} \rightarrow 5^2P_{\frac{3}{2}}$ to cool and confine the atoms. Thus our lasers will need to have a linewidth no larger than the linewidth of the p-state, which in this case is around 6 MHz.

3.2 Specs for transport

Atoms will need to be transported from the MOT to the center of the cavity, a distance of $\sim 4cm$. Ideally, we want a 2D plane of atoms in order to simplify the dynamics of the system. In practice will use an elliptical focus to transport a very thin cigar-shaped slice of the atoms. Once in the cavity, atoms will be excited to Rydberg states, which will grow the orbital radius of the atoms to the order of $1\mu m$. Thus we will want the minor-axis waist of the ellipse to have a radius $\sim 5\mu m$. Using an aspect ration of 40:1 gives an approximate plane

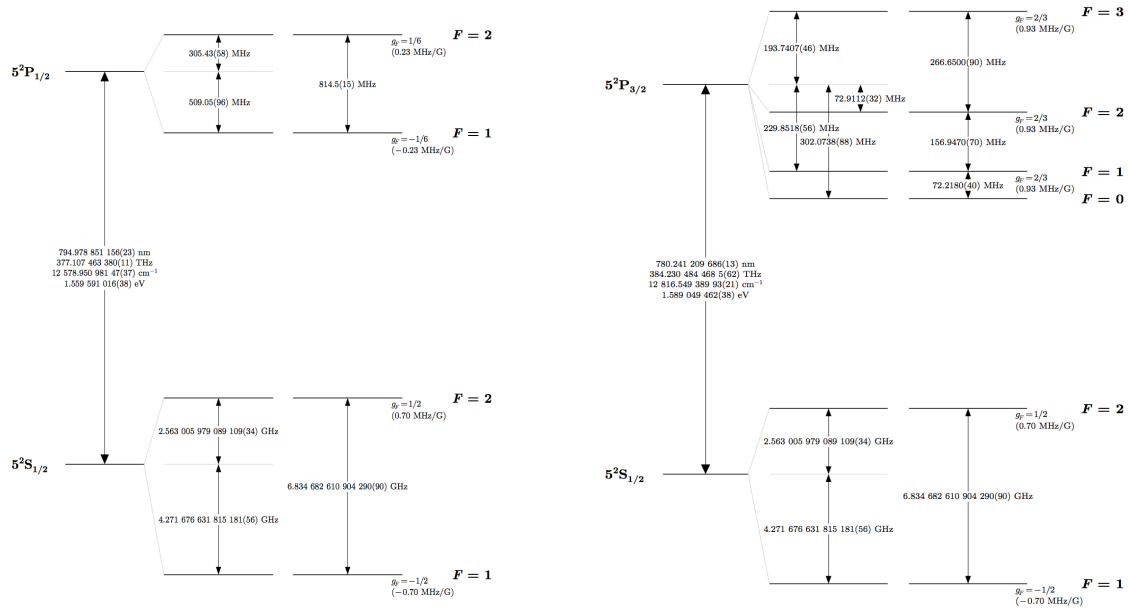

Figure 9: Atoms trapped in a MOT in the middle of a UHV chamber are transported via a standing wave with a moving focus into an optical cavity.

of atoms, so we use a major-axis waist radius of 200 μm. These numbers thus constrain the size of the focus of the standing wave. (Jon, what does this mean in terms of laser resolution?)

In order to facilitate numerous trials, the translation must be quick. We will want to be able to transport the atoms $4cm$ in $10ms$, requiring a transport speed of $4m/s$. To see how this speed requirement dictates the laser resolution we carry out the following calculation. The incoming field from one laser can be defined by $e^{ikx}e^{i2\pi\nu t}$ while the field from the counter-propagating laser is defined by $e^{-ikx}e^{i2\pi(\nu+\delta)t}$, where $k = \frac{2\pi}{\lambda}$ is the wave number, λ is the wavelength, ν is the frequency, and δ is the detuning of the lasers from one another. The total intensity of the two beams $I = |e^{ikx}e^{i2\pi\nu t} + e^{-ikx}e^{i2\pi(\nu+\delta)t}|^2 = 2(1 + cos(2kx - 2\pi\delta t))$. If we look at a point on the intensity curve, say where $2kx - 2\pi\delta t = 0$

we see that its position coordinate varies as

$$x = \frac{\pi\delta t}{k} = \frac{\delta\lambda t}{2} \tag{9}$$

Taking the time derivative we find the velocity

$$v = \frac{\delta\lambda}{2} \tag{10}$$

Equation[10] dictates how quickly the atoms in this conveyor belt will move, which we need to be $4m/s$. The lasers will have an off-resonant wavelength of $\lambda \approx 785nm$. Plugging both these numbers into Equation[10] gives the result $\delta = 10MHz$. (Jon, what does this demand of laser resolution?.) However, we can't simply switch δ from 0 to 10 back to 0 as extreme acceleration will jolt the atoms and cause heating. Instead, we need to smoothly increase δ from 0 to its maximum value and then smoothly decrease it back to 0. At the beginning of the transport process the acceleration is instantly switched from 0 to a and similarly from a to $-a$ after half the transportation distance. Shrader and Meschede (7) studied the acceleration as a function of transport efficiency and found that most transports with an acceleration $a < a_{max}$ had near 100% efficiency, where a_{max} is dependent on the potential depth.

3.3 Specs for EIT

Once the ground state Rb atoms have been transported into the cavity, we use a $480nm$ Rydberg EIT dressing beam and send in $780nm$ photons to create the Rydberg polaritons.

Precise frequencies, fast jumps and well-defined ramps between frequencies, and large ranges of achievable frequencies are among the requirements of our lasers. To achieve all of the above, we use an FPGA to control the frequency of

up to four lasers, which allows for time resolution of a couple hundred nanoseconds and frequencies ranging over 2 GHz.

4 Experimental Control Hardware

4.1 FPGA

An $FPGA$ is an array of an enormous number of configurable logic gates and memory blocks with programmable interconnects that allow the user to implement reconfigurable digital circuits, as well as input/output (IO) blocks that allow the $FPGA$ to communicate with the outside world. It is in this sense that the $FPGA$ programmer is programming the hardware. While microprocessors execute code on top of an operating system and application software which can lead to uncontrolled performance times, $FPGA$'s provide hardware-timed speed and reliability. Furthermore, the $FPGA$'s programmable logic units run in parallel, while the microprocessor executes a set of instructions sequentially, giving the $FPGA$ a major speed advantage through its massive parallel computing ability.

Our experiment will need the $FPGA$ to update the frequencies of multiple lasers, either by jumping or ramping to the next frequency, according to predefined schedules, which in general are different for each laser. As such, sequential logic will be necessary to sequentially update the frequency of each laser, looping over all lasers continuously. The best way to implement sequential logic on an $FPGA$ is with the mathematical construct of a finite state machine. This is an abstract machine that can be in any of a finite number of states. It can only be in one state at a time, but can move between different states depending on a number of triggers. This concept is easily mapped to the $FPGA$, where the machine will exist in one state and upon every clock cycle will either remain in

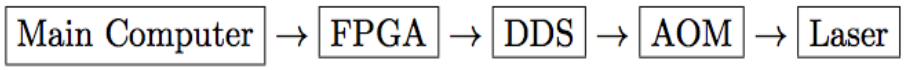

Figure 10: Laser frequency schedule information originates in the main experimental computer, and is sequentially passed to the FPGA, to the DDS, to the AOM, to the laser.

that state or switch to a new state based on the current state and the values of a few variables.

The $FPGA$ I developed for this function loads the frequency schedules for each laser from the main experiment computer onto its block RAM. For each schedule, the $FPGA$ outputs a sequence of 32-bit frequency tuning words to a direct digital synthesizer (DDS), which outputs a digital signal at a frequency proportional to the tuning words. This digital signal is then fed to an acousto-optic modulator (AOM), which will precisely tune a laser. The information flow is depicted in Figure[10]. The basic idea is that every clock cycle, the $FPGA$ should update one of the DDS's frequencies one step forward in its schedule, sequentially updating each DDS continuously in a loop.

The frequency sequences are each defined by a set of events. Each event includes a start time, a target frequency, and what I term a "step-size". The start time is in units of update cycle number, where the update cycle number increases every three $FPGA$ system clock cycles (as will be discussed below, this is because the state machine needs to visit three states per update, and each state is visited for one clock cycle each). The step-size is defined to be zero if the DDS should jump to the target frequency at the start time, otherwise, the DDS should increment its frequency by the step-size on each frequency update, until it reaches the target frequency. Since the $FPGA$ sequentially updates the four DDS's continuously, updating one DDS per update cycle count, the start times for, say, DDS_0 must coincide with times where

$(update\ cycle\ number)\%(number\ of\ \mathrm{DDS}'s) = 0$, while start times for DDS_1 must coincide with times where $(update\ cycle\ number)\%(number\ of\ \mathrm{DDS}'s) = 1$, and so on. Once all the start times have been adjusted, the list of events for all four DDS's is merged and sorted in order of start time. This list of events is then passed to the $FPGA$ over USB from the main computer and stored in the $FPGA$'s block RAM.

The finite state machine I implemented has some initial states that are only visited once, which load the first event from block RAM, set initial values for variables, and reset the DDS's. The machine is then sent to an idle state where is waits to receive a start trigger. Once triggered, the machine spends the vast majority of its time in three main states: the frequency modifying state, the IO update state, and the event loading state. In the frequency modifying state, the machine checks to see if the update cycle number is equal to the next event start time. If so, the new target frequency and step-size for the current DDS will be loaded into arrays that hold the current value for these variables for each DDS. If the step-size is zero, the new output frequency will be set to the target frequency, otherwise the new output frequency will be set to the current frequency for that DDS plus the new step-size. If the next event's start time has not yet been reached, the output frequency is simply updated based on the values stored in variable arrays for the current frequency and step-size of the DDS. The trick to controlling four DDS's at once is that while this output frequency tuning word is sent to all DDS's, only the DDS that receives an IO update trigger will update to the new frequency. Thus the next clock cycle brings the machine to the IO update state, which sends the IO trigger to the current DDS that needs to be updated. Finally, the machine is sent to the event loading state. If a new event occurred in the modify frequency state, the $FPGA$ will load the next event from the block RAM. Updating of the variable

array holding the current frequency for each DDS also occurs in this state to prepare for the modify frequency state. Thus, the $FPGA$ actually updates one DDS every three clock cycles, since it must go through three states per DDS update. As there are four DDS's, each is updated every 12 $FPGA$ clock cycles. The $FPGA$ has a 60 MHz clock, which means each DDS, and hence each laser, gets updated every 200 nanoseconds. Thus, the $FPGA$ provides us with the necessary fast and reliable time-control of multiple laser frequencies.

4.2 AOM

Each DDS is associated with one laser, and is able to tune that laser's frequency indirectly with the use of an AOM. The AOM functions by sending acoustic waves down the length of a crystal at a frequency set by a Piezo-electric transducer that squeezes the crystal at a frequency determined by the DDS output. The phonons create moving periodic planes of expansion and compression in the crystal. Since the index of refraction is dependent on density, it will be periodically modulated in the crystal. The incoming light thus essentially sees a diffraction grating as it moves through the crystal and will be deflected into collimated modes, similar to Bragg diffraction. However, this differs from Bragg diffraction in that the light is scattered off of a moving grating, which causes a frequency shift in the laser light, which works out to be equal to the acoustic wave frequency. Thus, changing the driving frequency of the AOM will change the frequency shift of the laser beam. It is in this manner that the DDS output is able to indirectly tune the laser. The DDS has a frequency resolution in the picohertz and this is how fine frequency resolution of the lasers is obtained. Together with the $FPGA$, the DDS's provide the extremely precise time- and frequency-control of multiple lasers, necessary to create a cold Rydberg gas capable of mediating controlled photon-photon interactions.

5 References

References

[1] R. W. Boyd, *Nonlinear optics*. Academic press (2003).

[2] M. Fleischhauer, A. Imamoglu, J. P. Marangos, *Reviews of Modern Physics* **77**(2), 633 (2005).

[3] K. Bergmann, H. Theuer, B. Shore, *Reviews of Modern Physics* **70**(3), 1003 (1998).

[4] H. Tanji-Suzuki, W. Chen, R. Landig, J. Simon, V. Vuletić, *Science* **333**(6047), 1266 (2011).

[5] E. Raab, M. Prentiss, A. Cable, S. Chu, D. E. Pritchard, *Physical Review Letters* **59**(23), 2631 (1987).

[6] C. Cohen-Tannoudji, W. D. Phillips, *Phys. Today* **43**(10), 33 (1990).

[7] D. Schrader, et al., *Applied Physics B* **73**(8), 819 (2001).