Lamb Shift in Atomic Hydrogen

1 *Background*

> *The spectrum of the hydrogen atom has proved to be the Rosetta stone of modern physics: once this pattern of lines had been deciphered much else could be understood.*
>
> —*Arthur Schawlow*

The goal in this experiment is to measure the Lamb shift in the Balmer α transition of atomic hydrogen. The Lamb shift cannot be explained by the Schrödinger or Dirac formulations of quantum mechanics. It can, however, be explained by a theory known as quantum electrodynamics—a theory whose development was intimately linked to experimental observation of the Lamb shift.

Note: Reading the article from *Scientific American,* "The spectrum of atomic hydrogen," by Hänsch, Schawlow and Series [1] (through p. 107) is strongly recommended. A copy is available for students to download from the class website. Although we do not use a dye laser, that section is still relevant reading as our tunable diode laser serves the same function as the dye laser described in the article.

Also, this write-up presents a number of exercises to the reader. As part of this experiment you are required to complete all exercises. Most can be done before doing the lab.

1.1 Classical description of the hydrogen atom

The spectrum of the hydrogen atom was the first to be described quantitatively and modeled from first principles. In 1885 Balmer discovered that the wavelengths of the then known lines in the hydrogen spectrum were described by the formula

$$\lambda = 3646 \left(\frac{n^2}{n^2 - 4} \right) \text{Å} \quad (1\text{Å} = 10^{-10}\text{m}) \tag{1}$$

where n is an integer and takes on the values $3, 4, 5, \ldots$ This series of lines is now known as the Balmer series, and it is the line corresponding to $n = 3$, known as the Balmer α line, that we will investigate in this experiment. As an aside, atomic transitions were first studied in the optical (humanly visible) portion of the electromagnetic spectrum. Transitions were observed using an instrument call a spectrograph which separates the incoming light by wavelength. Transitions at different wavelengths appear as distinct vertical lines in the spectrograph view field, hence the term "line" to denote a transition at a particular wavelength.

Exercise 1 *Use Balmer's formula to calculate the wavelength for the Balmer α line.*

In 1890 Rydberg discovered a more general form of Balmer's formula which, when applied to the hydrogen spectrum has the form

$$\frac{1}{\lambda} = R_H \left(\frac{1}{m^2} - \frac{1}{n^2} \right) \tag{2}$$

where R_H is a constant, known as the Rydberg constant (R takes on slightly different values depending on the mass of the atom under investigation), and m and n are integers with $m < n$. In

1913 Bohr proposed a model of the atom in which the (orbital) angular momentum of the electron is quantized in units of $\hbar$. Using this model he derived an expression for the energy levels E_n of the hydrogen atom

$$E_n = -\left(\frac{mZ^2e^4}{(4\pi\varepsilon_0)^2 2\hbar^2}\right)\frac{1}{n^2} \tag{3}$$

where m is the mass of electron, Z is the number of protons in the nucleus ($Z = 1$ for hydrogen), e is the charge of the electron and $n\hbar$ is the orbital angular momentum, with $n = 1, 2, 3, \ldots$ A transition from a state of higher energy E_{n_i} to one of lower energy E_{n_f} should result in the emission of radiation with energy

$$E = E_{n_i} - E_{n_f} = \left(\frac{mZ^2e^4}{(4\pi\varepsilon_0)^2 2\hbar^2}\right)\left[\frac{1}{n_f^2} - \frac{1}{n_i^2}\right] \tag{4}$$

where $n_i\hbar$ and $n_f\hbar$ are the electron orbital angular momenta of the initial and final states, respectively. Comparing this formula with the Rydberg formula and using the relation for radiated energy, $E = h\nu = hc/\lambda$, we find that R_H should be

$$R_H = \frac{me^4}{8\varepsilon_0^2 h^3 c}\ . \tag{5}$$

Refined spectroscopic measurements indicated that this expression was not quite correct—a situation that Bohr remedied by replacing the electron mass with the reduced mass μ of the atom,

$$\mu = \frac{mM}{M+m} \tag{6}$$

where M is the mass of the nucleus.

But other problems remained. In 1887 Michelson, using his interferometer to investigate the shape of spectral lines, discovered that the Balmer α line consisted of not a single line but of two lines separated by a fraction of an angstrom. Bohr's simple model of the atom had no means of accounting for this feature.

In 1916 Arnold Sommerfeld presented a model of the hydrogen atom allowing for the possibility of elliptical in addition to strictly circular electron orbits. Applying relativistic corrections to these elliptical orbits, Sommerfeld's model predicted the Balmer α line to consist of more than one component, just as Michelson had observed. But, being an *ad hoc* combination of classical and quantum physics, the Sommerfeld model was restricted in scope and left much to be desired. Quantitatively it was unable to account for features such as magnetic effects or the different intensities of the components of the Balmer α line.

1.2 Quantum mechanical description of the hydrogen atom

A more fundamental treatment of the hydrogen spectrum, as well as the spectra of other atoms, awaited the development of quantum mechanics in the 1920's. In 1928 Dirac presented a relativistic quantum theory with predicted energy levels for the hydrogen atom [2, pp. 284–6]:

$$E_{nj} = E_n\left[1 + \left(\frac{\alpha}{n}\right)^2\left(\frac{n}{j+1/2} - \frac{3}{4}\right)\right] \tag{7}$$

where

$$E_n = -\left(\frac{\mu e^4}{(4\pi\varepsilon_0)^2 2\hbar^2}\right)\frac{1}{n^2} \, . \tag{8}$$

The integer n is known as the principal quantum number, and as before can take on the values $n = 1, 2, 3, \ldots$ The α in Eq. (7) is a dimensionless constant known as the "fine-structure constant"; it represents the scale by which the energy corrections are reduced relative to the main energy term E_n that is due to the Coulomb (electrostatic) interaction between the proton and electron. The fine structure constant is a combination of various fundamental constants:

$$\alpha = \frac{e^2}{4\pi\varepsilon_0 \hbar c} \approx \frac{1}{137} \, . \tag{9}$$

The number j in Eq. (7) represents the total angular momentum of the electron; j is a combination of the orbital angular momentum and the spin angular momentum, and can take on the values $j = 1/2, 3/2, \ldots, n - 1/2$. Orbital angular momentum is specified by the integer l, which can take on the values $0, 1, \ldots, n-1$. Spin angular momentum is specified by the number s; for a single electron, $s = 1/2$. The total angular momentum number j can take on the values $l \pm s$ ($l \pm 1/2$ for a single electron) depending on whether the angular momenta are parallel or antiparallel. States with the same n but different values of j are called "sublevels."

An aside about notation of atomic states: states are specified using a long established notation, $n^{2s+1}X_j$, where n is the principle quantum number, s is the spin quantum number, X denotes the orbital angular momentum number (S for $l = 0$, P for $l = 1$, D for $l = 2$, F for $l = 3$), and j denotes the total angular momentum quantum number. A state with $n = 3, s = 1/2, l = 2$ and $j = 3/2$ is denoted as $3^2D_{3/2}$. For atomic hydrogen s is always 1/2 and so the $2s + 1$ term is often dropped from the specification, $3^2D_{3/2}$ becoming $3D_{3/2}$, for example.

The Dirac formulation accounts for energy shifts due to relativistic corrections and magnetic interactions between the orbital and spin magnetic moments. This equation tells us that the energy levels depend only on the principal quantum number n and the total angular momentum j, i.e., states (for a given n) with the same value of j will have the same energy. For example, a state with $l = 1$ and $j = 1/2$ (orbital and spin angular momenta antiparallel), designated $P_{1/2}$, should have the same energy as a state with $l = 0$ (no orbital angular momentum), designated $S_{1/2}$, and $j = s = 1/2$.

When applied to the Balmer α transition (the transition under investigation in this experiment) where $n_i = 3$ and $n_f = 2$, the Dirac equation predicts seven allowed transitions between the various sublevels of the $n = 3$ and $n = 2$ levels. Selection rules, which govern how much the angular momentum can change in a transition, prohibit transitions between certain sublevels. The selection rules applying to the type of transitions (electric dipole) under consideration here are: $\Delta l = \pm 1$, $\Delta j = 0, \pm 1$.

Exercise 2 *Use the selection rules and the notation described above to list all seven possible transitions between the $n = 2$ and $n = 3$ states in atomic hydrogen. Which transitions does the Dirac theory predict to have the same energy?*

Figure 1 shows the seven transitions, with the height of the solid lines in the lower part of the figure representing the strength of the transition. The curve above the seven lines represents the Doppler broadened profile that is observed when looking at the Balmer α radiation coming from a

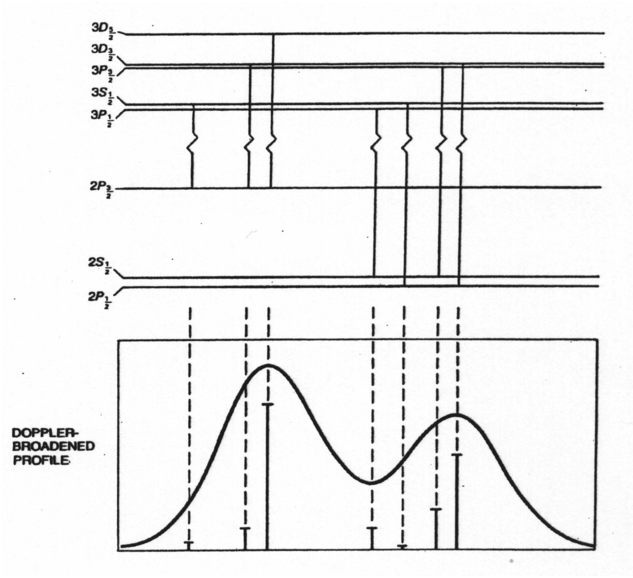

Figure 1: Doppler broadened profile of Balmer α lines.

hydrogen discharge lamp. The Doppler principle states that light emitted by an atom moving with velocity v relative to the observer, will, according to the observer, have frequency

$$\nu = \nu_0 \left(1 \pm v/c\right) \tag{10}$$

where ν_0 is the frequency of the light emitted by an atom at rest relative to the observer and c is the speed of light. The + sign is associated with the atom moving toward the observed, and the − sign with motion away from the observer.

Doppler broadening is due to the thermal motion of the atoms. The full width half maximum (FWHM) due to Doppler broadening is given by the formula

$$\Delta\nu = \frac{2\nu_0}{c}\sqrt{2k_B T \ln 2/M} \tag{11}$$

where k_B is Boltzmann's constant, T is the absolute temperature, M is the mass of the atom, and ν_0 is the center frequency of the transition [3, p. 189]. In the absence of Doppler broadening, the transitions of interest in this experiment have linewidths $\sim$ 100 MHz, considerably smaller than the Doppler broadened linewidth. (Linewidths observed in the absence of Doppler broadening and other broadening mechanisms, e.g. pressure broadening, are called natural linewidths, and are due to the finite lifetime of the excited state. The natural linewidth is determined by the uncertainty principle, $\Delta E \Delta t \geq \hbar$, where Δt is the excited state lifetime).

Note: Throughout this write-up energy is specified in terms of frequency, which is obtained by dividing the actual energy by Planck's constant, i.e., $\nu = E/h$.

Exercise 3 *Calculate the FWHM due to Doppler broadening for the hydrogen Balmer α transition at 300 K, and compare it to the energy difference between the $3D_{3/2}$ and $3P_{1/2}$ states.*

1.3 Experimental results that pointed the way

At the time Dirac presented his theory, experimenters had observed only two distinct Balmer α lines. Doppler broadening was clearly a limiting factor in attempting to resolve additional lines. What are the options for addressing this problem? The FWHM formula tells us that decreasing the temperature or increasing the mass will reduce the FWHM. In 1938 R. Williams exploited both of these factors by measuring the Balmer α spectrum from a deuterium discharge lamp cooled to 100 K [4]. He succeeded in resolving for the first time a third Balmer α component, the component corresponding to the $3P_{1/2} \rightarrow 2S_{1/2}$ transition. Figure 2 shows a plot of Williams' results for the Balmer α lines of both deuterium (D_α) and hydrogen (H_α). (In this case, energy units are in cm^{-1}, or wavelengths per cm. The value in cm^{-1} is often called the wavenumber. The conversion between energy stated in wavenumber and energy stated in frequency is 1 cm^{-1} = 29.979 GHz. To find the actual energy, multiply the energy stated in cm^{-1} by hc.)

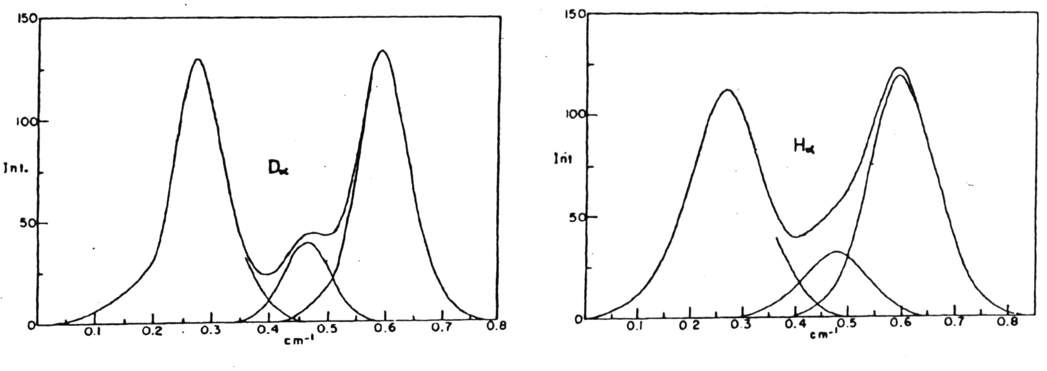

Figure 2: Intensity curves for D_α and H_α lines from the experiment of Williams (1938).

Exercise 4 *Calculate the FWHM due to Doppler broadening for the Balmer α transition in deuterium at 100 K, and compare it to the FWHM you calculated above.*

The Dirac theory predicted the separation between the $3D_{3/2} \rightarrow 2P_{1/2}$ and $3P_{1/2} \rightarrow 2S_{1/2}$ transitions to be just the energy difference between the $3D_{3/2}$ and $3P_{1/2}$ states (remember, the $2S_{1/2}$ and $2P_{1/2}$ states are predicted to have the same energy since they have the same value of j). The predicted separation is 3.24 GHz (the $3D_{3/2}$ state lies above $3P_{1/2}$ state).

But Williams' measurements showed this separation to be considerably larger. In Williams' words: "The most striking aspect of the data . . . is the large percent deviation from theory in the interval from component 2 to component 3." [4] (Component 2 is the $3D_{3/2} \rightarrow 2P_{1/2}$ transition and component 3 is the $3P_{1/2} \rightarrow 2S_{1/2}$ transition.) How could this discrepancy be explained? Was it possible that the $2S_{1/2}$ and $2P_{1/2}$ states were not degenerate? (Degeneracy means having the same energy.)

Exercise 5 *From the graph of Williams' deuterium data in Fig. 2, estimate the energy difference between the $3D_{3/2} \rightarrow 2P_{1/2}$ and $3P_{1/2} \rightarrow 2S_{1/2}$ transitions.*

In 1947 W. E. Lamb and R. C. Retherford performed a brilliant experiment in which they directly measured the energy difference between the $2S_{1/2}$ and $2P_{1/2}$ states [5]. In their experiment (see

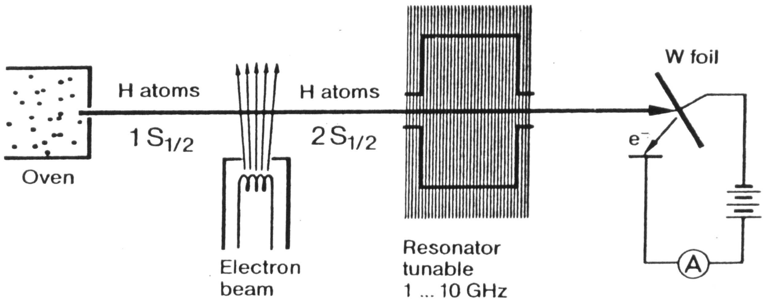

Figure 3: Configuration of the Lamb-Retherford experiment (1947).

Fig. 3), they prepared a beam of hydrogen atoms in the $1S_{1/2}$ ground state. The beam is bombarded with electrons with the result that some of the atoms are excited to the $2S_{1/2}$ state. Optical transitions from this state that go directly to the $1S_{1/2}$ ground state are prohibited by selection rules (which one?) resulting in a very long (~ 0.1 second) lifetime for the $2S_{1/2}$ state [3, p. 230]. The atoms then pass through a region where they are exposed to electromagnetic radiation with energy equal to the energy difference between the $2S_{1/2}$ and $2P_{1/2}$ states. The radiation has the effect of inducing transitions between these states, and once an atom is in the $2P_{1/2}$ state it decays to the ground state with a lifetime of 10^{-9} seconds.

After passing through the region with electromagnetic radiation, the atoms strike a tungsten foil. Upon striking the foil, atoms still in the $2S_{1/2}$ state decay to the ground state and in so doing liberate electrons from the foil in a process called Auger emission. By measuring the emission current from the foil with the electromagnetic radiation turned on (fewer electrons liberated) and with the radiation turned off (more electrons liberated), Lamb and Retherford were able to determine the energy difference between the $2S_{1/2}$ and $2P_{1/2}$ states. Their initial measurements indicated this difference to be about 1000 MHz.

Exercise 6 *Why was Doppler broadening not a problem in the measurements of Lamb and Retherford?*

1.4 QED picture of the hydrogen atom

The importance of the Lamb-Retherford measurement cannot be overstated—it was a milestone in experimental physics which spurred major developments in theoretical physics. In Dirac's own words: "No progress was made for 20 years. Then a development came, initiated by Lamb's discovery and explanation of the Lamb shift, which fundamentally changed the character of theoretical physics."

The Lamb-Retherford measurement led to the development of a theory known as quantum electrodynamics (QED). This theory proposes that the electromagnetic field consists of a collection of quantized harmonic oscillators, with the presence of a photon of energy $h\nu$ corresponding to the oscillator with frequency ν being in the first excited state. (A harmonic oscillator is one with a parabolic potential energy well.) A fundamental property of the quantized harmonic oscillator is that the energy in the ground state is not zero, but $h\nu/2$ [2, pp. 222–3]. This energy is called the zero-point energy. One way of thinking about the zero-point energy is that it is a requirement of the uncertainty principle: a particle in a potential well with zero kinetic energy would have both a well defined position (bottom of the well) and a well defined momentum ($p = 0$). This is forbidden by the uncertainty principle which requires that $\Delta x \Delta p \geq \hbar$.

Even in the absence of any external radiation, non-zero electromagnetic fields are present due to the fields associated with zero-point energy. These fields are known as "vacuum fluctuations," so named because they exist even in the absence, or "vacuum" of observable photons. The electromagnetic fields associated with vacuum fluctuations act on the electron to smear out its position relative to the nucleus. (The proton with its greater mass is less affected by vacuum fluctuations.) Because of this smearing, the electron charge is distributed over a larger volume of space, and less of its charge experiences the strongest Coulomb potential closest to the nucleus. The result is a reduction in the binding energy of the electron to the nucleus. Electrons in S states ($l = 0$) have a greater probability of being close to the nucleus than those in P or D states and so experience a greater reduction in binding energy than is the case for P or D states. The major effect of the vacuum fluctuations is to shift S states upward in energy relative to P states. QED calculations show that the contribution to the energy difference between the $2S_{1/2}$ and $2P_{1/2}$ states due to vacuum fluctuations is +1086 MHz. Calculations not based on formal QED theory, but still in fairly good quantitative agreement with the Lamb-Retherford measurement, were made shortly after experimental results were first announced. The calculation by Welton is straightforward and the interested reader can pursue this reference to get a feel for a simple theoretical treatment of the Lamb shift [6].

Exercise 7 *Explain how the upward shift in S state energy relative to a P state energy accounts for the energy difference observed by Williams (between the* $3D_{3/2} \rightarrow 2P_{1/2}$ *and* $3P_{1/2} \rightarrow 2S_{1/2}$ *transitions) being larger than predicted by the Dirac theory.*

A second QED phenomenon contributing to the energy shift between S states and P states is known as vacuum polarization. Virtual pairs of electrons and positrons are constantly created and annihilated in the Coulomb field of the proton. Being of opposite charge, virtual electrons are attracted to the proton, while positrons are repelled. The result is that the charge of the proton is partially screened by the virtual electrons. Electrons (ones bound to protons) penetrating the cloud of virtual particles will experience a stronger attraction to the nucleus than predicted by the Coulomb ($1/r$) potential, and hence will be more tightly bound to the nucleus. Electrons in S states ($l = 0$) have a greater probability of being close to the nucleus than those in P or D states and so experience a greater increase in binding energy than is the case for P or D states. QED calculations show that the contribution to the energy difference between the $2S_{1/2}$ and $2P_{1/2}$ states due to vacuum polarization is -27 MHz. Note that this energy shift is in opposite direction from that due to vacuum fluctuations. Other smaller effects such as reduced mass and finite size of the proton also contribute to the Lamb shift. The sum of these smaller contributions is approximately -1 MHz. Recent calculations for all effects contributing to the Lamb shift result in a value

$$1057.833\ (4)\ \text{MHz}\ [7].$$

The best experimental value for the $2S_{1/2} - 2P_{1/2}$ energy difference is

$$1057.845\ (3)\ \text{MHz}\ [7].$$

That these values are so close is an outstanding triumph of both experimental and theoretical physics.

1.5 Some final details

The astute observer will have noticed that the energy level diagram at the top of Fig. 1 already shows the splitting of the $2S_{1/2}$ and $2P_{1/2}$ states due to the Lamb shift. Indeed, it is this splitting that results in 7 distinct lines in the Balmer α transition instead of the 5 predicted by the Dirac equation. But what about the other splittings shown in the diagram—the splitting between the $3S_{1/2}$ and $3P_{1/2}$ states and that between the $3P_{3/2}$ and $3D_{3/2}$ states? What is the origin of these energy differences? It turns out these splittings are also manifestations of the Lamb shift. The same mechanisms, primarily vacuum fluctuations and vacuum polarization, that account for the splitting of the $2S_{1/2}$ and $2P_{1/2}$ states result in the splittings of the $n = 3$ states having the same value of total angular momentum j.

Exercise 8 *The magnitude of the Lamb shift goes as* n^{-3}. *Given a Lamb shift of 1058 MHz between the* $2S_{1/2}$ *and* $2P_{1/2}$ *states, predict the spitting between the* $3S_{1/2}$ *and* $3P_{1/2}$ *states. Compare your prediction to the measured value of 315 MHz.*

Figure 4 is a repeat of Fig. 1, but in addition shows an example of data similar to that you will acquire in this experiment. Careful scrutiny of this figure raises the following issue: the Lamb shift is defined as the energy separation between the $2S_{1/2}$ and $2P_{1/2}$ states. To measure this energy difference we need to compare two transitions where the initial states have the same energy and the final states are the $2S_{1/2}$ and $2P_{1/2}$ states. The transitions shown in Fig. 4 do not appear to satisfy this requirement. You will be measuring the energy difference between the $3D_{3/2} \rightarrow 2P_{1/2}$ transition and the $3P_{3/2} \rightarrow 2S_{1/2}$ transition, as shown in the figure. What about the energy difference between the $3D_{3/2}$ and the $3P_{3/2}$ states? This difference is 5.5 MHz—a non negligible amount relative to the precision of Lamb shift measurements [8].

It turns out that in the hydrogen discharge tube the $3D_{3/2}$ and $3P_{3/2}$ states are broadened and mixed by the linear Stark effect. This effect is a shift of atomic energy levels in the presence of an electric field [3, pp 219–227]. Atoms in the discharge tube experience not only the applied DC electric field (created by the high voltage power supply), but also local time-varying fields due to the plasma nature of the gas discharge. These strong time-varying fields smear out the Stark energy shifts thereby broadening the energy levels. When the broadened energy levels begin to overlap the two states can readily mix, resulting in a breakdown of the $\Delta l = \pm 1$ selection rule. The broadening and mixing result in these two states effectively sharing a common energy level, thus resolving the apparent discrepancy [9].

2 *The Experiment*

2.1 Saturation Spectroscopy

An inspection of Fig. 4 invites the question: how do you eliminate the Doppler broadening and get the resolution shown at the bottom of the figure—resolution that is essential to observe the Lamb shift with optical transitions? The answer: Doppler-free saturation spectroscopy. In this technique light with very narrow linewidth (several MHz for the diode laser used in this experiment) passes through a collection of hydrogen atoms undergoing Doppler-broadened transitions in a discharge tube. With its much smaller linewidth, the laser light excites (from the $n = 2$ state to the $n = 3$ state) only those atoms having just the right velocity to shift them into resonance with the laser light.

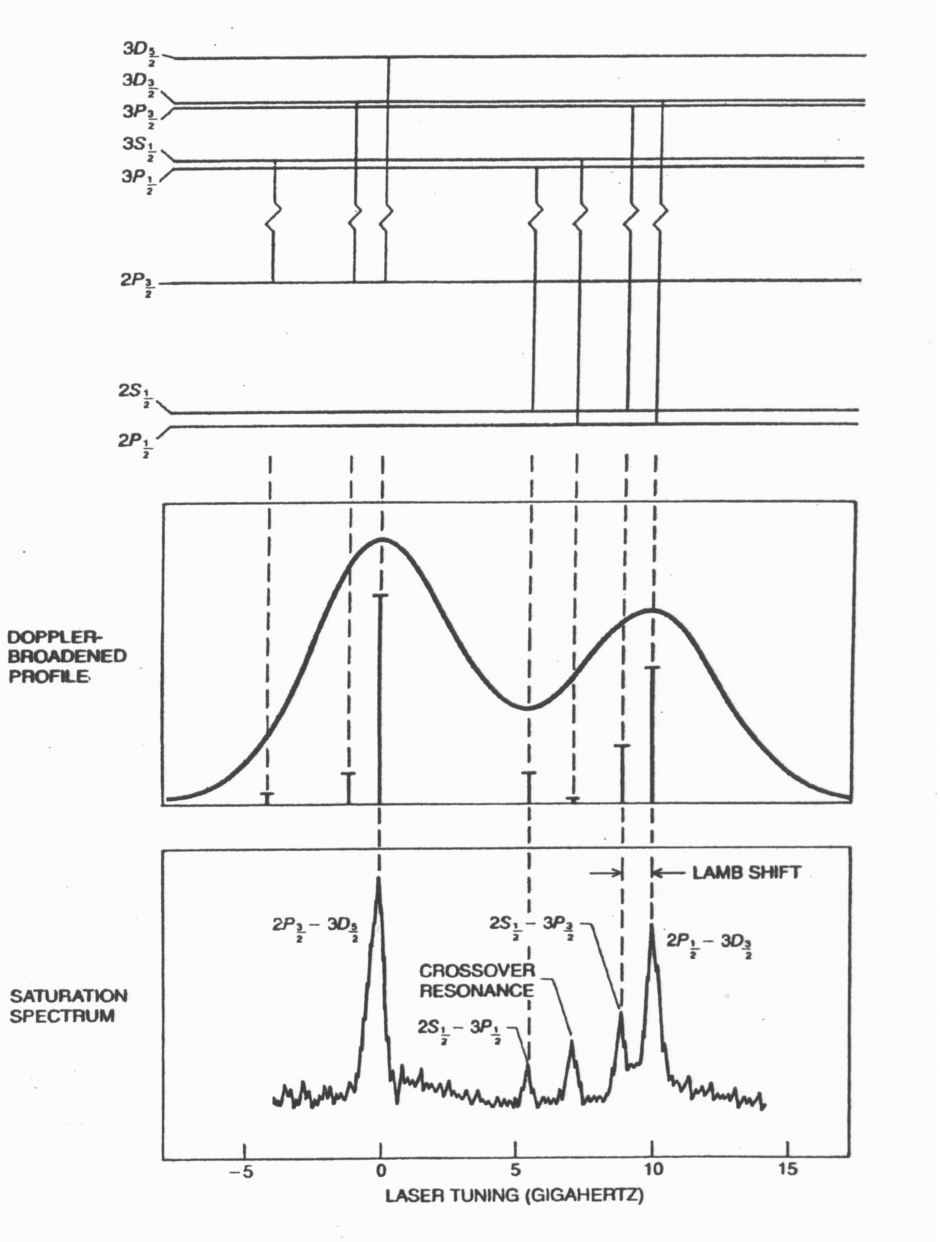

Figure 4: Top: Doppler broadened profile of Balmer α lines. Bottom: Balmer α lines resolved using saturation spectroscopy.

Our discharge tube is somewhat like the tube in a neon sign, only containing hydrogen instead of neon. When it is running, the tube in the experiment contains a mixture of molecules, neutral atoms, ions (in the form of protons) and free electrons.

The frequency (or wavelength) of the laser light can be varied continuously over the entire Doppler broadened transition. (Indeed, it was the development around 1970 of narrow linewidth, tunable lasers that made this kind of spectroscopy possible.) If the laser frequency ν_{laser} is set such that $\nu_{\text{laser}} < \nu_0$, the center frequency of a particular transition, an atom will come into resonance (i.e., will be excited from a lower energy state to a higher energy state) only if it is moving toward the light with velocity

$$v = \left(\frac{\nu_0 - \nu_{\text{laser}}}{\nu_0}\right) c \; . \tag{12}$$

Similarly, if $\nu_{\text{laser}} > \nu_0$ an atom will come into resonance only if it is moving away from the light with velocity

$$v = \left(\frac{\nu_{\text{laser}} - \nu_0}{\nu_0}\right) c \; . \tag{13}$$

Now consider a situation where two beams of laser light are propagating in *opposite* directions through the atoms in the discharge tube. One of the beams has considerably more power than the other and is called the pump beam or saturating beam. The less powerful beam is called the probe beam. The pump beam has sufficient power to promote a significant fraction of the atoms in its path from the lower, $n = 2$ level to the upper $n = 3$ level, and if it is powerful enough, it will equalize the $n = 2$ and $n = 3$ populations, thus "saturating" the transition.

A simplified schematic of the experiment is shown in Fig. 5. Note how the two beams are shown to overlap (with a highly exaggerated angle of crossing) inside the hydrogen discharge tube. In our experiment, the two beams are more nearly parallel than the beams in Fig. 5. We shall designate the axis along which the beams are propagating the x axis, as shown in the figure.

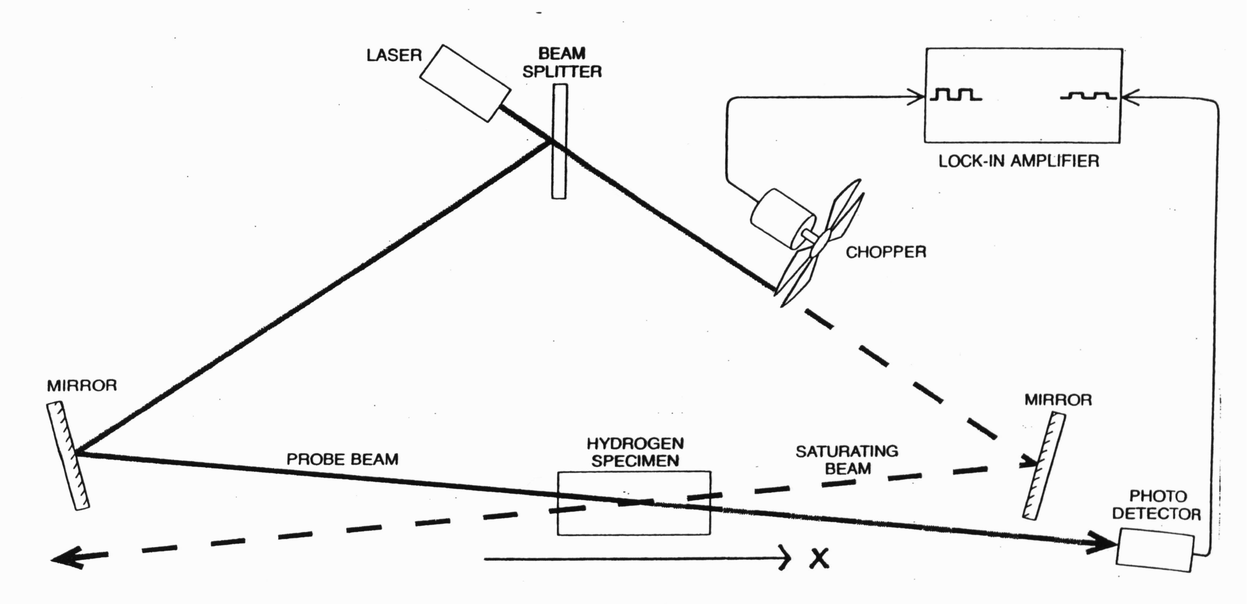

Figure 5: Experimental arrangement for saturation spectroscopy. The layout of optical components in our experiment is more complex; you should understand the function of the additional items.

The broad bell-shaped curves in Fig. 6 depict the distribution of velocities along the x axis of the hydrogen atoms. The distribution is due to the kinetic energy of the hydrogen atoms, and is directly related to the temperature of the gas. Note that the average velocity is 0—the atoms aren't going anywhere as a group. Those atoms with negative v_x are moving toward the probe beam and away from the pump beam; those with positive v_x are moving toward the pump beam and away from the probe beam. The negative dips in the bell curves represent the atoms which absorb the laser light at the frequency indicated; their number in turn depends on the intensity of the light and the fraction of the total which can be excited at that frequency.

The situation for $\nu_{\text{laser}} < \nu_0$ is shown in Fig. 6(a) and corresponds to the condition given in Eq. (12). The probe beam interacts with atoms moving toward it, and the pump beam interacts with atoms moving (in the opposite direction) toward it; since the laser frequency (in the laboratory reference frame) is less than the resonance frequency of an atom at rest, a given atom can only interact with that laser beam which is "blue-shifted" to the resonance. Thus, each beam interacts with a different set of atoms. The opposite situation is shown in Fig. 6(c), corresponding to Eq. (13), where an atom can only interact with a beam which is "red-shifted" towards the resonance.

Consider what would happen if only the probe beam in Fig. 5 were present. The detector measures the light coming through the hydrogen cell. When some of the light is absorbed by the hydrogen, the detector signal will drop. When one scans the laser frequency through the resonance, the amount of absorption will mirror the bell curves shown in Fig. 6. The absorption is the *convolution* of the laser intensity and the velocity distribution as a function of frequency. In other words, the signal shows the Doppler-broadened lineshapes depicted in Fig. 4.

When the pump beam is turned on, the bell-shaped distribution of atoms which can absorb light from the probe beam acquires an anomaly: the pump beam excites a portion of the atoms from the $n = 2$ to the $n = 3$ state, and as long as they stay in that state, they cannot absorb light from the probe beam. One might say that the pump beam "burns a hole" in the velocity distribution. If the laser frequency is such that the hole burned by the pump beam overlaps the range of velocities in which atoms can absorb light from the probe beam, the result will be a slight *decrease* in the light absorbed from the probe beam. One will see a slight uptick in the detector signal when this happens. Because the two beams are directed oppositely, this situation will occur only with those atoms where $v_x = 0$; this is shown in Fig. 6(b). An example of the detector output which shows the hole-burning effect is shown in Fig. 7. In the figure the upper trace is the detector signal which shows an overall decrease as the laser frequency is swept through the transition along with the small upward blips indicating hole burning.

The hole burning is seen only over the non-Doppler broadened linewidth of the transition centered at ν_0. When the laser frequency is outside this linewidth, the two beams are again interacting with different sets of atoms. As the frequency of the laser is swept through the non-Doppler broadened linewidth centered about ν_0, the probe beam intensity maps out just this linewidth—hence the term "Doppler-free saturation spectroscopy" to describe this technique. The trace at the bottom of Fig. 4 shows a plot of the probe beam intensity as the laser frequency is swept over the entire Doppler broadened Balmer α transition.

Exercise 9 *What are the relative populations of the $n = 2$ and $n = 3$ levels when they have been saturated by the pump beam? Answer: they are equal. Why is this?*

(a) $\nu_{laser} < \nu_0$

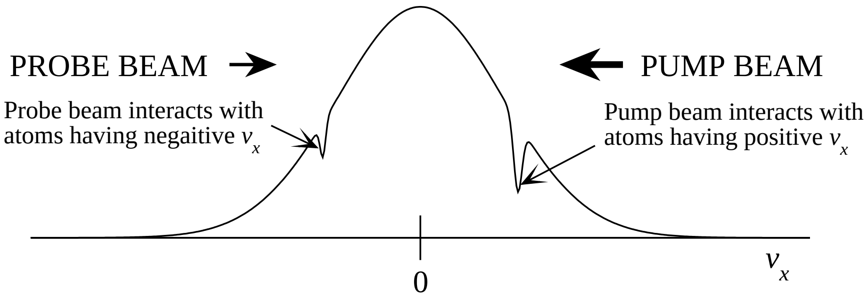

(b) $\nu_{laser} = \nu_0$

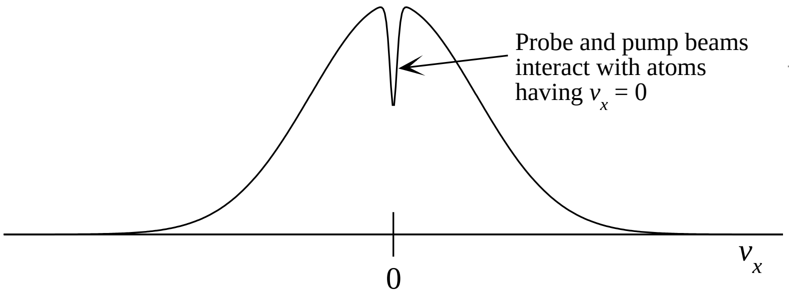

(c) $\nu_{laser} > \nu_0$

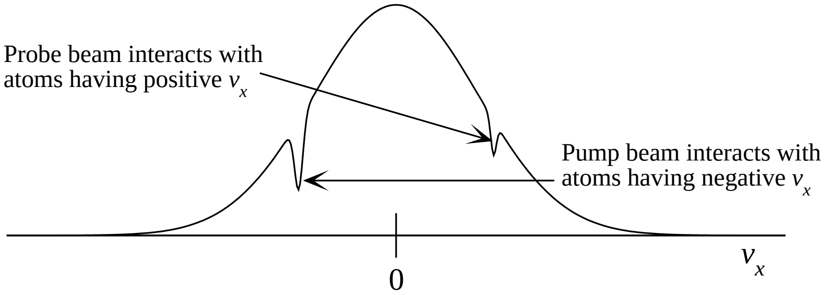

Figure 6: Absorbtion of pump and probe beams by a distribution of atoms as the laser frequency is swept through the resonance occuring at ν_0. The Bell-shaped curves show the usual Maxwellian velocity distribution along the x axis.

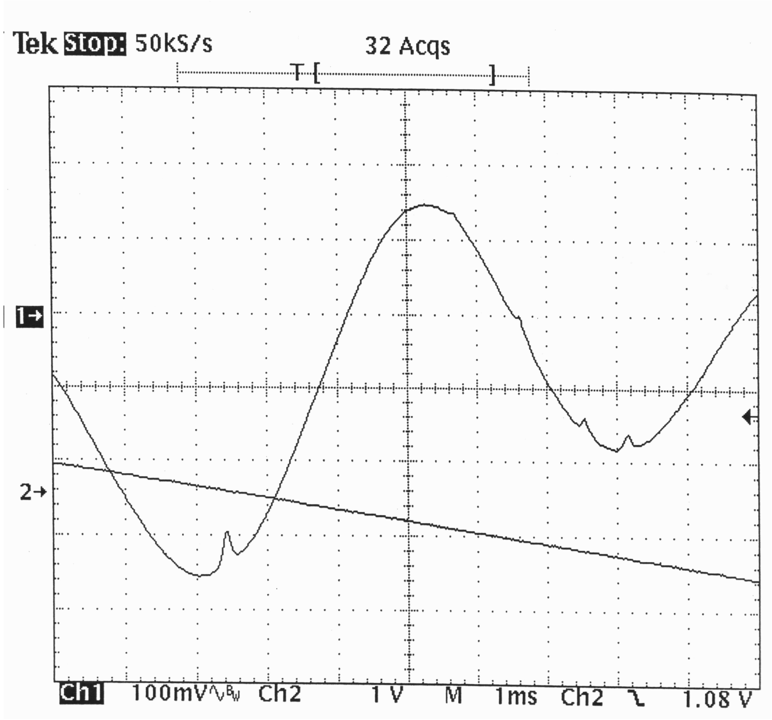

Figure 7: Oscilloscope trace showing Doppler broadened Balmer α absorbtion line and effect of hole burning. (Lower trace [straight line] is a control signal used to vary the laser frequency)

2.2 Lock-in detection

Figure 5 shows a fan-like device called a chopper. Where does this fit into the scheme of things? To answer this, first look at Fig. 7. This shows the intensity of the probe beam as the laser frequency is swept through the Doppler broadened transition. The curve is upside down relative to that shown in Fig. 4—we are looking at the remaining light after part of the beam has been absorbed by the hydrogen in the discharge tube; increasing intensity is in the upward direction.

Note the small blips in the Doppler broadened curve. These blips are due to hole burning: when the two beams interact with the same set of atoms, more of the probe beam intensity gets through. Note that there are 5 blips which appear to correspond to the 5 transitions shown in Fig. 4. At this point, you may want to claim that the measurement is over—we've located the various lines using the saturation spectroscopy technique. In a sense you are right, but the signal is small, and it would be difficult to precisely locate the various lines given the small size of the blips in Fig. 7. How can we do better? Enter the lock-in amplifier. When the pump beam is turned off there is no hole burning and all we get with the probe beam is the Doppler broadened profile. By turning the pump beam on and off periodically, we can make the blips appear and disappear periodically—the perfect situation for detection with a lock-in amplifier. The pump beam is turned on and off as it is alternately passed and blocked by the chopper blades at a frequency ν_{chop}, and a signal at this frequency is also fed from the chopper controller (not pictured in Fig. 5) to the lock-in reference input. *The lock-in is sensitive only to signals that appear and disappear at the reference (chopper) frequency; all it detects are the blips not the Doppler broadened curve.* By plotting the lock-in output as the laser frequency is slowly swept through the Doppler broadened linewidth, a plot resolving the various lines of the Balmer α transition will result. This is what is actually shown at the bottom of Fig. 4.

2.3 Crossover resonances

The Doppler-free saturation spectroscopy technique is clearly a very powerful tool for resolving closely spaced lines. It does, however, sometimes introduce extraneous features that are not part of the spectrum under investigation. Fig. 4 shows a resonance peak that looks like it might be associated with the $3S_{1/2} \rightarrow 2P_{1/2}$ transition, but is labeled "crossover resonance." What is this? Note that this peak is located halfway between $3P_{3/2} \rightarrow 2S_{1/2}$ and $3P_{1/2} \rightarrow 2S_{1/2}$ transitions, which share a common lower energy level. When the laser frequency is halfway between these transitions, the pump beam saturates the higher frequency transition ($3P_{3/2} \leftrightarrow 2S_{1/2}$) in atoms moving toward the pump beam and saturates the lower frequency transition ($3P_{1/2} \leftrightarrow 2S_{1/2}$) in atoms moving away from the pump beam. The probe beam also interacts with these same sets of atoms, but because they are saturated by the pump beam, the atoms absorb less of the probe beam resulting in the appearance of the crossover resonance. This resonance does not affect the measurements in this experiment, but it does have the effect of masking the weak $3S_{1/2} \rightarrow 2P_{1/2}$ transition.

2.4 Measuring the frequency shift

The standard method of measuring the frequency shift as the laser frequency is being swept is to send some of the laser light into a Fabry-Perot interferometer which will pass light at well defined frequency intervals. The way this works is as follows. A Fabry-perot interferometer is made of two partially silvered mirrors facing each other and aligned so that their faces are exactly parallel. The mirrors reflect the light between their two inner surfaces many times. If the wavelength of the light is such that an integer number of half-wavelengths fit in the space between the mirrors, then the multiple reflections allow a *standing wave* to be present. In other words, the mirrors form a resonant cavity for light waves, analogous to the sort of resonances one has in an organ pipe or flute for sound waves. Because the surfaces are only partially silvered, some of the light penetrates the mirror surface and escapes out of the cavity and can be picked up by a light detector. The detected light intensity is largest when the resonance condition is met. If laser light sent through the mirrors is swept through a range of frequencies, the detector will record an intensity oscillation corresponding to the change in frequency.

Let's put some mathematics to this description to see how the mirrors allow us to measure the frequency shift. Let the distance between the mirror surfaces be d, and let the laser wavelength be λ. A standing-wave resonance occurs when

$$n\frac{\lambda}{2} = d \; , \tag{14}$$

where n is some integer. For light waves in a vacuum, the frequency ν and wavelength are related by the well-known formula $\lambda\nu = c$, where c is the speed of light. Thus the frequency of this particular standing wave can be written

$$\nu = n\frac{c}{2d} \; . \tag{15}$$

If the frequency is increased, the wavelength will decrease, and the resonance condition will be lost, leading to a reduction in the detected light intensity. At some higher frequency, the next resonance will come into play, as there will be $n + 1$ half wavelengths between the two mirrors, and the light signal will peak again. If we call these two resonant frequencies ν_1 and ν_2, the difference between them is

$$\nu_2 - \nu_1 = (n+1)\frac{c}{2d} - n\frac{c}{2d} \, , \tag{16}$$

or

$$\Delta\nu = \frac{c}{2d} \, . \tag{17}$$

This expression is called the "free spectral range" of the Fabry-Perot interferometer, and it gives the amount of frequency shift between two peaks in a detector output sweep. Note that since d is in the denominator of the right side of Eq. (17), as we increase the distance between the mirrors, we are able to resolve ever smaller frequency shifts. This means we can make spectral measurements of very high resolution.

In our setup, the chopped pump beam is "picked off" by a mirror after it exits the discharge tube and sent through a Fabry-Perot intereferometer set at a fixed mirror spacing of about 10 centimeters. A detector is placed on the other side of the interferometer to measure the light intensity at the center of the beam. At this mirror spacing, the interferometer is highly sensitive to vibration, especially at the resonance peaks. So we use the lock-in trick on the intereferometer as well. With a long-enough lock-in time constant, the effects of vibration can be greatly reduced, and a scan of the output signal from the lock-in amplifier shows a relatively smooth variation as the laser frequency is swept.

You can compare the periodic variation in the Fabry-Perot signal to the peak locations in the hydrogen spectrum to calculate the frequency differences between various peaks. You will need to make a good measurement of the mirror separation in order to get the free spectral range. Details of how to do this and interpret your scan are given in the data analysis section of the experimental procedure.

2.5 Nobel Prizes

Willis Lamb was awarded the 1955 Nobel Prize in Physics "for his discoveries concerning the structure of the hydrogen spectrum."

Richard Feynman, Julian Schwinger and Sin-Itiro Tomonaga were awarded the 1965 Nobel Prize in Physics "for their fundamental work in quantum electrodynamics, with deep-ploughing consequences for the physics of fundamental particles."

Arthur Schawlow was awarded the 1981 Nobel Prize in Physics (along with Nicolaas Bloembergen) "for their contribution to the development of laser spectroscopy."

References

[1] Hänsch, T. W., A. L. Schawlow, and G. W. Series, "The spectrum of atomic hydrogen," *Scientific American,* March 1979, pp. 94–110.

[2] R. Eisberg and R. Resnick, *Quantum Physics of Atoms, Molecules, Solids, Nuclei and Particles, 2e,* Wiley, New York, (1985).

[3] B. Bransden and C. Joachain, *Physics of Atoms and Molecules,* Longman, London, (1983).

[4] Williams, R. C., "The fine structure of H_α and D_α under varying discharge conditions," *Phys. Rev.* **54**, 558 (1938).

[5] Lamb, W. E., and R. C. Retherford, "Fine structure of the hydrogen atom by a microwave method," *Phys. Rev.* **72**, 241 (1947).

[6] Welton, T. A., "Some observable effects of the quantum-mechanical fluctuations of the electromagnetic field," *Phys. Rev.* **74**, 1157 (1948).

[7] Eides, M. I., H. Grotch and V. A. Shelyuto, "Theory of light hydrogenlike atoms," *Physics Reports,* **342**, 63–261 (2001).

[8] Weber, E. W., and J. E. Goldsmith, "Double-quantum saturation spectroscopy in hydrogen: Measurement of the $3P_{3/2}$-$3D_{3/2}$ Lamb shift," *Phys. Rev. Lett.* **41**, 940 (1978).

[9] Hansch, T. W., I. S. Shahin, A. L. Schawlow, "Optical resolution of the Lamb shift in atomic hydrogen by laser saturation spectroscopy," *Nature* **235**, 63 (1972).

[10] Erickson, G. W., "Energy levels of one-electron atoms," *J. Phys. Chem. Ref. Data*, Vol. **6**, No. 3, p. 831 (1977).

3 Experimental Procedures

In this experiment, the Lamb shift is observed in the $n = 2$ to $n = 3$ (Balmer α) transitions in atomic hydrogen. Molecular hydrogen is generated by electrolysis of water, and then dissociated into atomic hydrogen inside a high voltage discharge tube. The discharge tube also excites the hydrogen into the higher n states so that the Balmer α line can be explored. Tunable, monochromatic laser light is directed through the tube from both ends, and by manipulating the counter-propagating beams appropriately, one can extract non-Doppler-broadened (Doppler-free) profiles of the various transitions, thereby resolving them and allowing for observation of the Lamb shift.

Caution: there are several hazards associated with the Lamb shift apparatus:

1. **High voltage and electrical shock.**
2. **High laser power (5+ milliwatts).** ***Never stare into direct or reflected beam!***
3. **Generation of hydrogen gas.**

It is important to follow the written guidelines, and to handle and adjust the equipment deliberately and safely. If you have any questions about how to do anything with the apparatus, consult the TA, Professor or Lab Manager.

3.1 The gas handling system and flow control

Note: As of Spring 2020, the hydrogen generator system which created hydrogen gas by electrolysis of a KOH salt solution has been replaced with a simpler, less hazardous arrangement that uses purified water vapor as the discharge medium.

It is necessary for the gas handling system to be thoroughly evacuated before allowing the purified water vapor into the discharge lamp. The TA, Professor or Lab Manager will start the vacuum system earlier in the day and after the appropriate period of time, start the high voltage discharge. **Do not attempt to start the high voltage discharge without assistance from the TA, professor or Lab Manager.**

In the discharge lamp, the water vapor molecules are disassociated by the discharge, and atomic hydrogen mainly fills the lamp cavity. The oxygen ions are attracted to the negative-potential electrode (cathode), and combine with the copper metal at that electrode to produce the dark coating seen on the glass. The positive-potential electrode (anode) attracts electrons and negative hydrogen ions. Thus, when the discharge is glowing, one end of the lamp is more purely pink than the other, indicating a higher density of hydrogen.

During the experiment water vapor flows into the discharge tube continuously. It is metered into one end of the tube through valve #3, and is pumped from the other end by a mechanical pump.

Before the the lamp can be lit, the system must be pumped out. Check that this is the case. The valves should be set as follows:

#0 Stem valve on glass flask with the long neck: CLOSED

#1 open (black handle vertical)

#2 closed (black handle horizontal)

#3 Flow control valve: open to approximately 0.25” (see below)

#4 open

#5 open

#6 open (This will be closed after the system is open to water vapor.)

#7 (valve on vacuum pump) open

The vapor flow rate and thereby the tube pressure are sensitive to the flow valve setting; changing the setting by one division (0.001” = one mark on the handle) may result in a pressure change of 10’s of 100 millitorr (pressure is sometimes stated in microns and denoted by the symbol μ; 1 μ = 1 millitorr). During operation of the discharge tube, the flow valve should set to create a steady pressure of about 350 millitorr on the “Hastings Vacuum Gauge” meter. Because the vapor pressure of water at room temperature is fairly low (about 20 torr, compared to atmospheric pressure of 760 torr) the valve needs to be open fairly wide: between 0.250” and 0.275”. One full turn of the micrometer knob opens the valve by 0.025” and each full turn of the knob reveals one more small increment on the barrel of the micrometer.

If the system is well pumped (below 30 millitorr) you can allow the water vapor into the system. There are two steps:

1. Close valve #6. This will force the vapor to pass through the metering valve #3 and into the lamp.
2. Carefully but firmly grasp the stem on the water flask near the top and open the stem valve (#0). Back the valve stem out by 5 mm, as seen on the scale in the valve. (The lower of the two O-rings should be at the 5 mm mark.) You should see the pressure start to rise. If the metering valve is set properly, the pressure should settle to about 350 millitorr.

If the tube pressure is not in the range of 300 millitorr, turn the knob on the metering valve (#3) clockwise to reduce the pressure or counter clockwise to increase it. Have patience! The change in the pressure measurement is slow. Turn the knob a little and wait for it to equilibrate. Adjust the flow control valve (#3) so that the discharge tube pressure is close to 350 millitorr.

3.2 Checking the laser beam alignment

A comment about the laser: This is an expensive ($15,000), high quality instrument. Please treat it respectfully and keep the operational parameters (laser current, laser power, sweep amplitude and frequency) within the recommended limits.

Check that the Laser Sweep switch on the switch box is in the “Off” position.

If it is not already power on, turn on the laser controller (key switch). The controller performs a self-check, and then displays the value of the various parameters. After several minutes, the diode temperature should stabilize at 20.0° C, the recommended operating temperature. Check that the diode current (far right knob) is set at or near zero. Turn on the laser power (button on upper left), and when the light on the button stops blinking, slowly turn up the diode current, noting the jump in beam intensity at about 39 mA as the diode changes from an LED into a laser. Increase the diode current until the indicated laser power is about 2.0 mW (current should be $\approx$ 45 mA). **Do not exceed 52 mA current (about 5 mW).**

With the laser on, check the alignment of the laser beams relative to each other and to the discharge tube. This is done by holding an index card up to the round holes in the wood box. First confirm that the mechanical chopper is off, and that the chopper blade is not blocking the pump beam. If the blade is blocking the pump beam, nudge it gently to a new position where it no longer blocks the beam. Standing on the "laser" side of the wood box, hold an index card up to the hole on the right side of the box. Take a second index card and block the pump beam near the chopper. Now only the probe beam will show up on the index card at the circular hole. When the pump beam is unblocked, it should appear on the card just to the side of and at the same level as the probe beam. It may partially overlap the spot from the pump beam. Now hold up the index card to the hole on the left side of the box. As before, block and unblock the pump beam. The pump beam should appear on the opposite side of the probe beam as on the right side of the box, again at the same level. When this geometry is obtained, the beams will cross inside the discharge tube—a necessary condition to observe the Doppler-free line profiles. The crossing point is chosen to be in the region of highest hydrogen density, near the anode (positive) end of the tube.

Both the pump and probe beams should pass cleanly through the discharge tube without hitting the walls. If the beams seem distorted or blurred out on the index cards, or if you can see that they are hitting the walls of the tube, consult with the TA, Professor or Lab Manager as to how to achieve the proper beam alignment. Please do not try to align it yourself. The alignment must follow a particular set of steps, and random knob turning will probably make it worse not better.

Finally, check that the probe beam falls on the detector properly. Block the probe beam with an index card. The detector current should decrease by approximately 10 μA.

Exercise 10 *Make a sketch of the optics and show the path of the two beams. Describe the effect of the various optical components; in particular, explain the role of the small mirror M4 and the iris. Also show the components used to measure the frequency shift.*

At this time, turn on the power to the remaining electronic instruments: function generators, both lock-in amplifiers (power switch on the SR830 lock-in is on the back of the unit), and the scope (if they are not already on).

3.3 The discharge tube

Confirm with the TA, Professor or Lab Manager that the vapor flow has been running for the necessary period of time and the system is ready to light. **Do not attempt to light the discharge tube without supervision; equipment damage or personal injury may result from improper use of high-voltage equipment.** When the tube is lit, confirm that the following experimental parameters are at or near their nominal values:

Discharge tube high voltage: +1900 – +2200 volts (start with +2000 V).
Discharge tube pressure: 200–400 millitorr.
Discharge tube current: 16–21 mA.

When all is well with the discharge, it will have a beautiful and relatively uniform magenta coloration. If it looks orange or white, impurities are present; such impurities will degrade the quality of the Lamb shift data. Running the tube for a while ($\approx 1/2$ hour) should flush out the impurities and result in the desired magenta coloration. If not, consult the TA, Professor or Lab Manager for advice. (Note, it is not possible to remove *all* of the impurities.)

It is not uncommon for the lamp to go out after it is started. The system needs to stabilize and the lamp should warm somewhat. If it goes out, check that the pressure has not dropped below 350 millitorr and start it again with the Tesla coil. You can cause the pressure to rise quickly by closing the vacuum-vent valve (#5) briefly. Please ask for help if this occurs, unless you have been trained to start the lamp.

If the lamp suddenly drops in intensity, but does not go out completely, close valve #5 for a moment and allow the pressure to go up. You should see it get suddenly brighter after the pressure has risen enough. Then open valve #5, and notice that the lamp gets brighter still.

In general, the lamp will stabilize and get more uniformly pink as it remains on. The white color at the cathode end will diminish over this time.

It is interesting to view the discharge tube through one of the small diffraction grating slides. If, when looking through a slide, you see three strongly colored lines—the Balmer α (red), β (turquoise) , and γ (violet) lines—with very little color between these, then the gas is fairly pure. With significant impurities, you will see other colors, mainly yellow.

3.4 Observing Doppler-broadened lines

The two function generators put out triangular waveforms that sweep the laser wavelength through the transitions of interest. The analog function generator puts out the fast sweep (about 40 Hz) that is used when adjusting the system. The digital generator (SRS DS345) puts out the slow sub-Hz sweep that is used when acquiring data with the X-Y recorder. When the Laser Sweep is on, the center wavelength of the sweep can be varied continuously over a small range with the Piezo Voltage control.

The most challenging part of operating this experiment is locating the Balmer-α lines. Practice is required because the coarse wavelength adjustment can be inconsistent. The nominal value for the red line is 656.28 nm, but the wavelength readout on the laser control can be anywhere from 656.1 to 656.3 nm. You may not be able to see the Doppler broadened absorption lines without first making a coarse adjustment of the laser wavelength, as described in the second paragraph below. To start your search, proceed as follows.

On the switchbox, set the “CH2” switch to “Sweep” and the “Sweep Speed” switch to “Fast”. Now look at the sweep signal on the scope. Check that the sweep frequency is approximately 40 Hz, and that the amplitude is about 3 volts peak-to-peak. (This is on the analog generator.) **The *absolute maximum* sweep voltage is 5 volts peak-to-peak.** Check that the “Scope Input” switch is set to the PDA-700 position. (The PDA-700 is the little amplifier that converts the photodiode output current to a voltage.) Flip the “Laser Sweep” switch to the “ON” position. If you are lucky, you will see something like the lower curve in Fig. 8, where the two humps are the Doppler-broadened absorption curves. (Beam intensity increases going upward; going through a resonance removes energy from the beam, thus reducing the detected intensity). If you see only a hint of these curves, vary the Piezo Voltage control to try to bring the curves onto the scope display. If you are unable to get a display of the broadened absorption curves, the laser wavelength will need to be coarsely adjusted.

To complicate your search, the laser intensity also varies during the wavelength scan with the piezo, so it may not be obvious when you have found the lines. The dips for the absorption are much narrower and deeper than this variation, but it takes some practice to spot them. Ask for help if you do not see them right away.

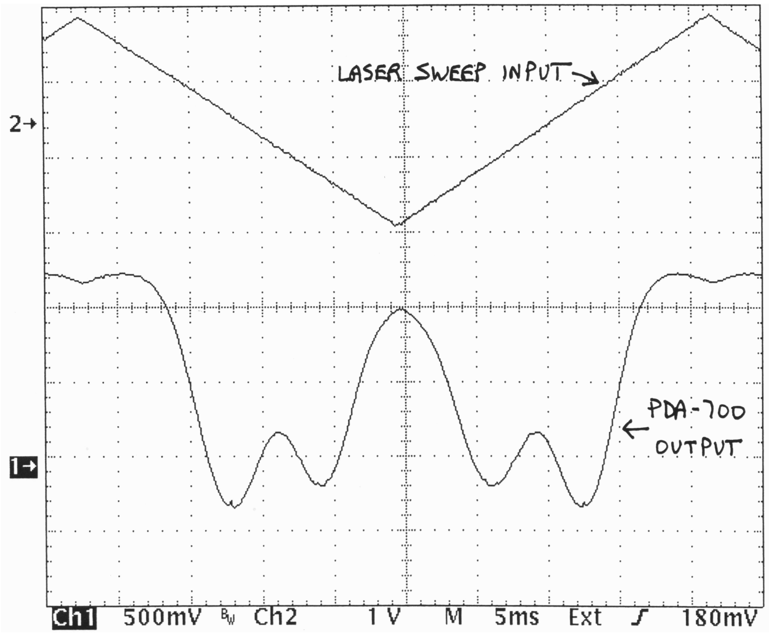

Figure 8: Oscilloscope trace showing Doppler broadened Balmer α absorption lines. Equipment settings: tube pressure = 190 μ, laser power ≈ 5 mW, sweep frequency = 23 Hz, PDA-700 gain = 20 μA.

To make coarse changes in the wavelength, flip the “Laser Sweep” switch to “OFF”, set the Piezo Voltage control to 50, and push the Track button above the Wavelength Adjust knob. Change the wavelength by turning the knob two clicks at a time. Note how the wavelength display shows resolution to 0.01 nm when the wavelength knob is turned, but only 0.1 nm otherwise. After changing the coarse wavelength, turn the Track button off, the “Laser Sweep” back on and hunt for the absorption humps by varying the Piezo Voltage control.

The laser sometimes exhibits an undesirable behavior known as *mode-hopping,* with a result shown in Fig. 9 Mode-hopping results when the cavity + diode combination jumps from one standing wave mode to the next; i.e., the number of half-wavelengths of light in the cavity changes by one. Mode hopping usually happens when the Track button and the laser sweep signal are on simultaneously. This condition is signaled by a persistent “beeping” from the laser head, as the picomotor tries to counteract the effect of the piezo. If the mode-hopping appears and persists after the coarse adjust is turned off, varying the Piezo Voltage, reducing (or increasing) the laser power, or *slightly* increasing the laser temperature control setting will usually cause the mode hopping to go away. ***Please ask for help if you think you need to increase the laser power or temperature.***

3.5 Locating and recording Doppler-free lines and frequency shift signal

With just one cycle of the sweep waveform and absorption curves displayed (as in Fig. 8), turn on the AVERAGE feature of the scope (see ACQUIRE menu) and set it to 16 or so sweeps/average. A close inspection of the absorption curves will reveal several small positive-going blips. At the location of these blips, the probe beam intensity increases (remember, increasing intensity is up on the scope). To see the blips more clearly, adjust the scope settings to obtain a picture similar to Fig. 7.

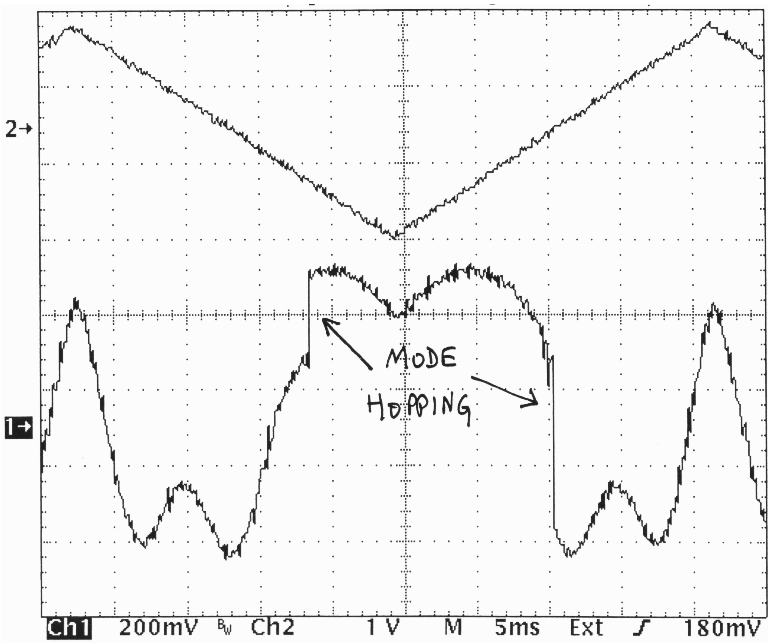

Figure 9: Oscilloscope trace showing mode hopping instability in tunable laser.

Block off the pump beam and observe what happens to the blips. Unblock the pump beam and again observe the effect. Why does the intensity of the probe beam increase at these locations when the pump beam is on and decrease when the pump beam is off? It is due to the "hole burning" action of the pump beam.

Obtain scope screen pictures of the absorption peaks with and without the "hole-burning" blips to include with your notebook.

When the mechanical chopper is on, it performs the same function as you did in blocking and unblocking the pump beam, i.e., it switches the Doppler-free hole burning effect on and off. A reference signal at the chopping frequency is connected to the Reference input of the lock-in, and the output of the PDA-700 amplifier is connected to the lock-in Signal input. In this configuration, the lock-in detects the Doppler-free hole burning only, since this is the signal that is being switched on and off at the lock-in reference (chopping) frequency.

Once you have obtained a scope display something like Fig. 7, you can proceed to record data with the XY recorder as follows. This sequence of steps is designed to help you first locate the peaks clearly, using the scope to view the lock-in amplifiers with somewhat fast time constants, and then to get the most out of the lock-in amps by running them with very slow time constants along with slow sweep speeds when you make the XY recordings.

1. Check that the digital function generator (DS345) settings are as follows:
    - Amplitude: 1.5 volts peak to peak (this is the voltage going into a 50 Ω load; the voltage will be twice this value going into a high impedance load).
    - Waveform: triangular.
    - Frequency: approximately 0.1 Hz.
2. On the Switchbox, flip the Sweep Speed from Fast to Slow.

3. You should see the signals on the oscilloscope going up and down slowly. To see the absorption dips again, slow the oscilloscope's Time/Div setting down, and use the horizontal position knob to put the display window to the beginning of the scope's memory. You should see the trace coming in from the right side of the screen and scrolling leftward. (Note the little bar indicator at the top of the scope display as you turn the position knob. The beginning of the memory corresponds to the right end of the bar.)

4. Once you can see the absorption dips scrolling across the screen, turn on the mechanical chopper. It may take a moment to start spinning, but when it does, the chopping/reference frequency will be displayed on the SR830 lock-in, and it should be set to around 300 Hz (the value is not critical). Try turning the big Frequency knob on the chopper control box up and down, and see (and hear) how the chopper frequency changes on the lock-in amplifiers' frequency displays.

5. Suggested initial settings for the probe-beam (SR830) lock-in are:

   Time constant = 30 mS

   12 dB/octave roll-off

   Signal Input: A, DC coupled, Grounded

   20 mV sensitivity

   High reserve

   DC coupled

   Line filter ON

   Channel One settings:
      Display: R
      Output: Display
      Ratio: Off
      Expand: Off
      Offset: Off

6. On the switch box, select the "Lock-in" for the "Scope Input". (Check that the "Laser Sweep" switch is on.) The output from the lock-in is bigger than the output from the PDA-700 amps, so your signal may be off-screen. Adjust the Volts/Div setting counterclockwise to comfortably show the peaks from the spectrum as the scan voltage varies. Have patience. You should be able to see a big peak followed by a space and then four smaller peaks (or the other way around, depending on the scan direction). The peaks may be rounded off.

7. On the switch box, select the “F-P Signal” for the “CH2” input. The initial settings on the EG&G 5210 lock-in (the one with the black face panel) should be as follows:

   Sensitivity

   Input: V, A
   Float/Ground switch: Out (grounded)
   Sens. setting: 300 mV

   Filters

   Mode: BP
   Set Freq: Track
   Line reject: “F” and “2F” both on

   Display 1

   Ref F: selected
   Reference: “Int” and “2F” both off

   Display 2

   Time constant 30 mS
   Slope: 12 dB
   Dyn Res: Norm
   Display: RΘ

8. Turn on the PDA-700 amplifier that is connected to the Fabry-Perot interferometer detector (if not already on). Set the “Range” switch to .xxxx μA. Cover the entrance of the detector pinhole with your finger and adjust the “Offset” knob so that the display reads zero (or as close as you can make it). Release your finger, and note the change in current when the laser light hits the detector.

9. Now you should be seeing two kinds of traces on the oscilloscope: A CH1 trace that shows the output of the probe-beam signal as processed by the SR830 lock-in amp, with the Doppler-free absorption peaks, and a CH2 trace that shows the output of the interferometer detector, with the periodic peaks showing the frequency shift of the laser.

   If you don’t see these traces clearly, try adjusting the scope controls and/or lock-in sensitivity levels. (At the time constant settings chosen, there will be a fair amount of noise.) If you still can’t get the traces to appear properly, turn off the chopper power (make sure that after the blade stops it is not blocking the beam), and switch the “Sweep speed” back to the fast sweep and verify that the blips due to hole burning are present on the Doppler-broadened lines displayed on the scope.

10. Once you have see the correct signals from both lock-in amps on the scope, reduce the sweep frequency on the digital generator (DS345) to something like 0.01 Hz. A slower frequency is necessary to allow the lock-in to track the changes in the signal as the laser output passes through the various resonances.

11. On the Switch Box, select the “Lock-in” for the “Plotter Y-axis” input and turn the “X-axis” and “Y-axis” switches for the “X-Y plotter” on.

12. Start up the computer-based LabVIEW Chart Recorder application on the computer with the National Instruments interface box attached to it. It should have three cables from the box connecting to outputs on the switchbox. (The cables may be hard to follow, as they snake underneath the wooden box with the discharge lamp.) The channels on the Chart

Recorder are connected to the switchbox as follows: Channel 1 = Y1-axis, Channel 2 = Y2-axis, Channel 3 = X-axis.

13. After the Chart Recorder application starts, click on the tabs at the top left to set the scan parameters. The following settings should produce a nicely visible set of traces (if the setting is not specified, the default value is acceptable):

    Channel 1 Check Enable Ch 1.

    Channel 2 Check Enable Ch 2. Under Ch 2 Offset set it to 4 or 5. This will put the Fabry-Perot signal at the top of the scan. You can change this later.

    Channel 3 Check Enable Ch 3, Invert Signal and Set as X-axis.

    Sample Clock The sample clock controls how often a measurement is recorded by the interface box. We will use a little signal averaging to reduce noise. Check Pre-bin samples and set Bins/Read-cycle to 10.

14. Go back to the Scan Controls tab, click TAKE SCAN, and see what you get. *You may need to click more than once. The Chart Recorder application is new as of Spring 2023, and has a number of small bugs when this set of instructions was written. It does collect accurate data, but its controls need some work.*

15. Time your scan stop and start so that you capture only one direction. Notice that if you just let it run, the data from two directions do not overlap exactly. This is because of hysteresis in the laser wavelength shift (a common issue with piezoelectric position controllers).

16. You can save a scan in local memory to look at it later by clicking RETAIN SCAN. If you give the scan a name (in Scan Name) it will save this name with the scan. Otherwise it will use the default. Subsequent RETAIN SCANs will save more scans with the assigned name and a number.

17. A scan you want to save permanently should be saved to disk by clicking SAVE TO FILE. It is recommended to use the file extension `.wfm`, which indicates that it is a LabVIEW "waveform" file (a binary format). You can also load old scan files into the application to study and manipulate them.

18. Try different scan rates and lock-in time constants by changing the frequency on the digital function generator (DS345). For example, set the frequency to 0.003 Hz, and then set the time constants on both lock-in amps to 300 ms. Try a scan at this slow speed. With some practice you can stop a scan, retain it and start a new one at the end of each direction.

19. When you are happy with the settings for your scans, record at least two for further analysis: one with the scan moving from left to right, and another with the scan moving from right to left. You will notice that these two directions do not line up! Notice the frequency shift trace. You may also see that the peak spacing is not uniform. This effect and how to deal with it is discussed in the analysis section. ***Do not touch anything on the optics table while your scans are running!*** **If you do, you will disturb the Fabry-Perot mirrors and cause a glitch in your frequency shift scan.**

20. Don't forget to RETAIN and SAVE TO FILE your best scans before stopping the application.

3.6 Shut Down

1. Turn down the high voltage setting to 0 volts. Flip the High Voltage switch to the Stand-by/Reset position. Turn off the high voltage power supply.
2. Close valves #0 (on the water flask) and #1 and open valve #6 so the system can pump down to vacuum.
3. Turn the laser diode current down to 0, turn off the laser power (push button), and turn off the laser controller (key).
4. Turn off the remaining electronic instruments.
5. After the discharge tube has pumped down to the 20–30 millitorr range, close valves #4 and #5, and unplug the vacuum pump.
6. Then open valve #2 to vent the pump input to atmosphere.
7. Cover up the optics.

4 Measure the mirror spacing

In order to relate the peaks in the interferometer scan to the frequency shift, you must know the free spectral range of the interferometer, and for this, you need to know the spacing d of the Fabry-Perot mirrors.

Use a caliper to measure the distance between the inside surfaces of the aluminum mirror mounts by placing the caliper between the mounting flanges at the edge. Be careful to make the caliper perpendicular to the flange surfaces. You should measure the distance between three different points and then average the results. If you don't understand these instructions, ask for help! Note the variation in the measurements and calculate an uncertainty.

Also, write down the amount of recess that each mirror surface has relative to the surface of the mount. These numbers are written on the sides of each aluminum mounting flange.

5 Data Analysis

To complete the data analysis, you should have two XY plots showing the peaks of the Balmer-α line—one for each scan direction along with the peaks from the Fabry-Perot detector, plus a good measurement of the Fabry-Perot mirror spacing.

First, load one of your good scan into the Chart Recorder application. Then use the "capture cursor" feature to record peak positions. Here is how: hover the mouse cursor over the cursor legend at the lower right of the chart. With the mouse cursor over the red "Measure" cursor, right-click to bring up a pop-up menu. From this, choose "Center Cursor". You should see a read cross-hair appear in the middle of the screen. With the mouse, grab (left-click and hold) the center of the cross-hairs. You should be able to move it freely over the plot.

Then position the cursor over the left-most peak in the Fabry-Perot scan, and then click `CAPTURE CURSOR`. The X and Y values of cursor will appear in the table. You may at this point type a label (i.e., “1”) in the first column. Proceed in this way to record and label all the peaks in the Fabry-Perot scan. The peaks from this are relatively broad, so you will need to use your best judgment to estimate their centers. Discuss any choices that might seem questionable in your report.

Click EXPORT TABLE to save the table to an Excel spreadsheet. (Give it a minute; the application as to start Excel unless it has already been opened previously.) Save the temporary Excel file to disk with a useful name.

Repeat the above process to locate and record the peaks in the hydrogen spectrum scan. Note: if you need to “zoom in” to the plot to control the cursor more accurately, you can use the Crop Visible tab under Plot Controls and Adjustments. First click SELECT X REGION and two blue cursor bars will appear on the screen. Move these to isolate the region you wish to examine more closely. Then click the button again and the positions will be retained (note the Low and High boxes). Select Enable and the data will be trimmed to the region of interest. If autoscaling of the axes is on, the data will be zoomed in. If you unclick Enable, the original data set will appear, and you can select another region.

Before proceeding further, compare the peak spacing between similar peaks for the two separate plots. In particular, note any regular trends in the peak-spacing of the interferometer scans. How much change is there from the beginning of the scan to the end? How much change is there in the region of the hydrogen peaks?

Your goal is to get the spacing of the peaks in frequency units, and you can see that the change in frequency is not uniform across the scan, but it is close. To make the calibration, one could fit a line to the Fabry-Perot peak positions and obtain the average spacing in units of the x-axis from the slope of the line of peak positions versus peak number. Applying this calibration to the hydrogen peaks should give their frequency in units of the free spectral range of the F-P set, relative to an arbitrary zero. A better approach that incorporates the nonlinearity in the scan is to fit the F-P peak positions versus peak number with a third-order polynomial that follows the smooth curve of the positions. This would make a calibration curve to apply to the hydrogen peak positions, from which one could obtain a “linearized” correction. Whatever method you decide to use, justify it in your report. Calculate the Lamb shift for both of your scans and compare them to each other. Do they agree within the uncertainty you would calculate? Average your two results and compare the average to the accepted value given in the theoretical discussion of the effect. Discuss.

The systematic variation in interferometer peak spacing is due to hysteresis in the piezo-electric response of the laser scanning mechanism. The piezo crystal flexes under the influence of an electric field. This flexing is what moves one of the mirrors in the laser cavity, which then changes the resonant frequency and hence the color of the laser. The hysteresis is a “memory” of the crystal’s previous state, e.g., if the crystal was in a compressed state, it tends to stay that way even if the applied voltage is changed so that it should expand. As a consequence, the crystal flexes less at the beginning of the scan than at the end, and the interferometer peaks are more widely spaced at the beginning than at the end. In high-precision work, the interferometer is output is coupled to the laser and to an independent frequency standard in order to counteract this hysteresis. Our setup does not have this added control feature, but you should be able to obtain the Lamb shift to within a few percent.

You should also calculate the separation between the peaks at the extreme ends of the spectrum, and the separation from the smallest peak in the middle to the largest peak at the end and compare your result to the accepted values:

9890 MHz from the $2P_{3/2} \rightarrow 3D_{5/2}$ transition to the $2P_{1/2} \rightarrow 3D_{3/2}$ transition.

5583 MHz from the $2P_{3/2} \rightarrow 3D_{5/2}$ transition to the $2S_{1/2} \rightarrow 3P_{1/2}$ transition.

These are derived by comparing published values for frequency intervals between the $3D$ states and between the $2P$ states [10].

Prepared by J. Stoltenberg, D. Pengra, and R. Van Dyck
`lamb_shift.tex -- Updated 21 April 2023`