Vol. 2

Chapter 2.4

Page 1:

Good morning, everyone. Welcome back to Phys 608.

I'm Distinguished Professor Dr M A Gondal, and in our last lecture, we established the foundations of saturation spectroscopy, a powerful method for overcoming the Doppler broadening that so often masks the true structure of atomic and molecular transitions.

Today, we are going to dive into a related, but in many ways superior, technique. As you can see from the title slide, we'll be covering Chapter 2, Section 4: Polarization Spectroscopy.

This method is not just an incremental improvement; it represents a significant leap in sensitivity and offers some wonderfully elegant features that make it an indispensable tool in the modern laser spectroscopy laboratory.

So, let’s begin.

Page 2:

So, let's start with the core motivation. Why do we need another Doppler-free technique? What is the fundamental idea behind "Polarization" Spectroscopy?

The central goal, as stated here, is to detect Doppler-free changes in the populations of *sub-levels*... and this is a key term we'll unpack shortly... by monitoring a change in the *polarization* of a probe beam, rather than simply a change in its transmitted intensity.

This immediately sets up a critical contrast with the saturation spectroscopy we've already discussed. Let's recall the principle of saturation spectroscopy. We use a strong pump beam to saturate a transition for a specific velocity class of atoms, effectively "bleaching" the sample for that group. A weak, counter-propagating probe beam then experiences reduced absorption when it is tuned to interact with that same velocity class. The signal we measure is that small reduction in absorption, or equivalently, a small *increase* in transmission. Therefore, the signal in saturation spectroscopy is proportional to the change in the absorption coefficient, which we denote as $\Delta \alpha$ $\Delta\alpha$.

The problem is that this small change, this $\Delta \alpha$ $\Delta\alpha$, is detected on top of the large, transmitted intensity of the probe beam itself. You are looking for a small bump on a large background. Any fluctuation or noise in the laser's intensity directly translates into noise in your signal, fundamentally limiting the achievable signal-to-noise ratio. Polarization spectroscopy offers a clever and elegant way to circumvent this very problem.

Page 3:

So, how does it accomplish this? As the first point here states, the polarization technique measures the *rotation* and/or the induced *ellipticity* of the probe beam's polarization state. This means our signal is proportional to simultaneous changes in both the absorption coefficient, Delta alpha, and, crucially, the refractive index, Delta en. The interplay between these two is the heart of the technique.

This different approach brings with it some immense practical advantages, which we will detail mathematically later, but let's introduce them conceptually now.

First, and most importantly, it offers a much higher intrinsic signal-to-noise ratio. This is because, in its ideal form, polarization spectroscopy is a "zero-background" technique. We set up our polarizers in a crossed configuration, so that without any interaction in the sample, *no* probe light reaches the detector. The signal is the small amount of light that *leaks* through the second polarizer because its polarization has been altered by the sample. So, instead of looking for a small change on a large background, we are looking for a small signal on a nearly zero background. This is a monumental advantage for achieving high sensitivity.

Second, the technique can be configured to directly produce a dispersion-shaped line profile. As we'll see, this kind of antisymmetric signal, with a steep, linear slope at the line center, is incredibly convenient for locking the frequency of a laser to an atomic transition. It provides a perfect error signal without needing any extra frequency modulation on the laser itself, which simplifies the experimental setup considerably.

Now for a bit of historical context...

Page 4:

This technique was truly brought to prominence in the mid-1970s by Carl Wieman and Theodor Hänsch, as noted here by the references.

Their work demonstrated its power and simplicity, and since then, polarization spectroscopy has become a standard, go-to tool for high-resolution, Doppler-free measurements. It proudly stands alongside saturated absorption and two-photon spectroscopy as one of the pillars of modern laser spectroscopy.

It's a technique that every student in this field should not only know about but deeply understand.

Page 5:

To make the contrast absolutely clear, let's look at these two block diagrams side-by-side.

On the left, in panel (a), we have the familiar setup for Saturation Spectroscopy. A laser beam is split. The weaker probe beam passes through the sample cell and onto a detector. The stronger pump beam is sent in the opposite direction, counter-propagating through the cell. The detector simply measures the total intensity of the probe beam. The signal, as we've discussed, is a small change in absorption, $\Delta\alpha$, appearing as a small peak on top of the transmitted probe beam's intensity.

Now, look to the right, at panel (b), Polarization Spectroscopy. The setup starts similarly, with a laser and a beam splitter creating a pump and probe. But notice the crucial differences. First, the probe beam passes through a linear polarizer, which we'll call P1, *before* it enters the sample cell. This prepares the probe in a well-defined state of linear polarization. Second, look at the pump beam. After the mirror, it passes through a quarter-wave plate, labeled $\lambda/4$. This converts the linearly polarized pump into circularly polarized light. This is key, as we will see. Finally, and most importantly, after the sample cell and *before* the detector, the probe beam must pass through a second polarizer, the analyzer, which we'll call P2. This analyzer is oriented to be "crossed" with the first polarizer, P1, meaning its transmission axis is rotated by 90 degrees relative to P1.

In this crossed configuration, if the sample had no effect, the linearly polarized probe from P1 would be completely blocked by P2, and the detector would see nothing. The signal arises only when the pump beam alters the sample, making it optically active, which in turn rotates or changes the probe's polarization, allowing some of it to leak through P2. The signal is therefore directly dependent on both the change in absorption, $\Delta\alpha$, and the change in refractive index, Δn, induced by the pump. This is a fundamentally different, and as we will prove, a far more sensitive detection scheme.

Page 6:

Alright, let’s break down the core experimental layout in more detail, building on the diagram we just saw.

First, the Optical Path Split. We begin with a single laser source. Critically, this laser must be both monochromatic—meaning it has a very narrow frequency distribution—and tunable, so we can scan its frequency across the atomic or molecular resonance. The output of this laser is divided by a beam splitter into two distinct paths.

The first path becomes Beam 1, which we call the *probe* beam. As its name implies, its job is to probe the state of the sample. It is essential that the probe beam has a weak intensity, which we'll denote as I_1. We need it to be weak so that it doesn't significantly alter the atomic populations itself. It should act as a passive observer.

The second path is Beam 2, the *pump* beam. Its job is to actively change the sample. Therefore, it must have a strong intensity, denoted I_2. This strong intensity is what allows us to saturate the transition for a specific group of atoms, which is the first step in creating the conditions for our signal.

Page 7: Polarization Elements

Next, let's consider the Polarization Elements, which are the heart of this technique.

The weak probe beam, on its way to the sample, passes through a high-quality linear polarizer, which we've labeled P 1 P_1. The function of P 1 P_1 is to define a clean, initial polarization axis for the probe. We can think of this as our reference.

The strong pump beam, on its path, passes through a quarter-wave plate, often denoted as λ / 4 $\lambda/4$. Assuming the pump is linearly polarized to begin with, and its axis is at 45 degrees to the wave plate's axes, this optical element converts the pump beam from linear to circular polarization. For our discussion, we'll assume it's converted to right-hand circular polarization, which in spectroscopic notation is represented by σ + σ^+. We'll see very shortly why having a circularly polarized pump is so important.

Now, for the Sample and Detection phase. Both the probe and the circularly polarized pump traverse the same vapor cell containing our atomic or molecular sample. And, just as in saturation spectroscopy, they do so in *opposite directions*. This counter-propagating geometry is absolutely essential for selecting the zero-velocity class of atoms and achieving a Doppler-free signal.

Page 8:

Finally, what happens after the probe beam has traversed the sample?

It exits the cell and immediately encounters a second linear polarizer, P 2 P_2, which we call the analyzer. As mentioned before, this analyzer is "crossed" with the first polarizer, P 1 P_1. This means its transmission axis is set at a 90 ∘ 90° angle to the initial polarization of the probe.

So, if nothing had happened to the probe's polarization inside the cell, it would be completely extinguished by P 2 P_2. The detector, labeled D D, sits right after the analyzer P 2 P_2. Its job is to measure any light that manages to "leak" through. This leaked light is our signal. The very existence of a signal tells us that the sample, under the influence of the pump beam, has altered the probe's polarization. Our entire measurement, then, consists of detecting this faint light against a dark background.

Page 9:

This slide provides a more detailed schematic of the core experimental layout, summarizing everything we've just discussed. Let's trace the paths one more time to solidify our understanding.

We start with the tunable laser on the far left. Its output hits a beam splitter.

Let's follow the probe beam first—the one that goes straight through. It passes through the initial polarizer, P 1 P_1, which defines its polarization as linear, let's say vertical for the sake of this diagram. It then travels through the sample cell, where it interacts with the atoms. After the cell, it encounters the crossed analyzer, P 2 P_2, which has a horizontal transmission axis in this picture. Finally, any light that gets through P 2 P_2 hits the detector, D D.

Now, let's trace the pump beam—the one that's reflected downwards by the beam splitter. It reflects off mirror M 1 M_1, then passes through a quarter-wave plate, which converts its polarization to circular, let's say σ + σ^+. It's then directed by mirror M 2 M_2 to enter the sample cell from the right, traveling in the opposite direction to the probe. Inside the cell, this strong, circularly polarized pump beam interacts with the atoms. As the text box explains, it saturates the transition and, as we'll see, induces an *anisotropy* in the sample.

This pump-induced anisotropy is what affects the probe beam. It causes the probe's plane of polarization to rotate slightly. As you can see in the diagram, the polarization vector of the probe, which was initially vertical, is now slightly tilted as it emerges from the cell. This small rotated component can now pass through the horizontal analyzer P 2 P_2, generating our signal at the detector. This entire diagram beautifully illustrates the cause-and-effect chain of the experiment.

Page 10:

Now we arrive at the physics at the heart of the technique. How, exactly, does the pump beam create this anisotropy in the sample? The answer lies in the Angular Momentum Selection Rules. This slide focuses on the pump's job.

Let's consider a generic electric-dipole transition in an atom or molecule. The transition is between a lower state and an upper state. Each state is characterized by a total angular momentum quantum number, which we'll call J. Let's denote the lower state by J'' and the upper state by J'.

Furthermore, these energy levels are degenerate. In the absence of external fields, the orientation of the angular momentum vector J in space is not specified. However, the propagation direction of our light beam provides a natural quantization axis, which we'll call the z-axis. The projection of the total angular momentum J onto this z-axis is given by the magnetic quantum number, M. So our transition is more completely described as going from a specific sub-level $|J'', M''\rangle$ to an upper sub-level $|J', M'\rangle$.

Now, here is the crucial point. The light itself carries angular momentum, and for a transition to occur, angular momentum must be conserved. For our circularly polarized pump beam—which we've assumed to be σ^+—the photon carries one unit of angular momentum along the direction of propagation. The strict selection rule for such a transition is that the magnetic quantum number M must change by $+1$. That is, ΔM, which is $M' - M''$, must equal $+1$.

Page 11:

What is the consequence of this strict selection rule, $\Delta M = +1$?

The first consequence is profound: Not every M'' sub-level in the lower state is addressable by our σ^+ pump beam. The pump cannot interact with all the atoms in the ground state equally.

Let's take a concrete example, as shown in point 2. Consider a P-branch transition, which by definition is a transition where $\Delta J = -1$. So, $J' = J'' - 1$.

Now, think about the sub-level in the lower state that has the maximum possible projection of angular momentum along the z-axis, which is $M'' = +J''$. For this atom to be excited by a σ^+ photon, it would need to transition to an upper state with $M' = M'' + 1$, which would be $J'' + 1$. But the upper state manifold, with total angular momentum J', only has M' sub-levels up to a maximum value of J', which is $J'' - 1$! There is no available state for the atom to go to. Therefore, the sub-level $M'' = +J''$ is completely immune to the pump beam. It cannot participate in a σ^+ transition.

The result of this selective interaction is an *unequal depletion* of the ground state sub-levels. The pump beam removes population from some M'' states, but not others. This creates a non-uniform, or *anisotropic*, sub-level population distribution. The gas is no longer isotropic; it now has a preferred orientation in space, imprinted upon it by the pump beam's polarization.

Page 12:

This diagram provides a perfect visual illustration of the selection rules at work. Let's analyze it carefully.

The title says we are looking at a P-branch transition, where $\Delta J = -1$, with σ^+ polarized light, which means $\Delta M = +1$. The specific example shown is a transition from a lower state with J ″ = 2 $J'' = 2$ to an upper state with J ′ = 1 $J' = 1$.

At the bottom, we see the energy sub-levels of the lower J ″ = 2 $J'' = 2$ state. The magnetic quantum number, M ″ M'', can take values from − J $-J$ to + J $+J$, so we have levels for M ″ = − 2 , − 1 , 0 , + 1 , $M'' = -2, -1, 0, +1,$ and + 2 $+2$.

At the top, we see the sub-levels for the upper J ′ = 1 $J' = 1$ state. Here, M ′ M' can be − 1 , 0 , $-1, 0,$ or + 1 $+1$.

Now, let's apply our σ + σ^+ selection rule: M ′ = M ″ + 1 $M' = M'' + 1$.

- An atom in the M ″ = − 2 $M'' = -2$ state can be pumped to the M ′ = − 1 $M' = -1$ state. This is allowed, and we see a blue arrow for this transition. - An atom in the M ″ = − 1 $M'' = -1$ state can be pumped to the M ′ = 0 $M' = 0$ state. This is also allowed, shown by the second arrow. - An atom in the M ″ = 0 $M'' = 0$ state can be pumped to the M ′ = + 1 $M' = +1$ state. This transition is also allowed, and is the third arrow shown.

Now, consider the M ″ = + 1 $M'' = +1$ state. A transition would require an M ′ = + 2 $M' = +2$ state. But look at the upper manifold! There is no M ′ = + 2 $M' = +2$ sub-level. So, this transition is forbidden.

Similarly, for the M ″ = + 2 $M'' = +2$ state, a transition would require an M ′ = + 3 $M' = +3$ state, which also does not exist.

The clear result is that the strong pump beam depletes the populations of the M ″ = − 2 , − 1 , $M'' = -2, -1,$ and 0 0 sub-levels, while leaving the populations of the M ″ = + 1 $M'' = +1$ and + 2 $+2$ sub-levels completely untouched. This is the very definition of creating an anisotropic population distribution.

Page 13:

So, we've established that the pump creates an unequal M M-population. What does this mean macroscopically? It means we've induced a preferred orientation of the molecular or atomic angular momentum, J J. Before the pump, these vectors pointed in all directions randomly. Now, there is a net alignment or orientation.

This leads to the crucial consequence: the medium itself becomes optically anisotropic. An isotropic medium responds the same way to light, regardless of the light's polarization or direction. Our pumped medium is no longer isotropic. Specifically, it will now respond differently to right-hand circularly polarized light, σ + σ^+, versus left-hand circularly polarized light, σ − σ^-.

This optical anisotropy manifests in two distinct, measurable ways.

First, we get a difference in the absorption coefficients for the two circular polarizations. The absorption coefficient for σ + σ^+ light, which we call α + α_+, will not be equal to the absorption coefficient for σ − σ^- light, α − α_-. This phenomenon is known as *circular dichroism*.

Page 14:

The second macroscopic consequence is a difference in the refractive indices. The refractive index experienced by σ + σ^+ light, n + n_+, will not be equal to the refractive index experienced by σ^- light, n_-. This phenomenon is known as *circular birefringence*, which is analogous to the birefringence you may have studied in crystals, but here it's induced in a gas by our pump beam.

Now, let's bring the probe beam back into the picture. Remember, our probe beam is linearly polarized. A key insight from classical optics is that any linearly polarized light can be mathematically described as a perfect, equal-amplitude superposition of right-hand (σ^+) and left-hand (σ^-) circularly polarized light.

So, when this linearly polarized probe enters our now-anisotropic medium, its two circular components are treated differently. The difference in refractive indices, $n_+ \neq n_-$, means one component travels slightly slower than the other. This introduces a relative phase shift between them. When they recombine upon exiting the sample, this phase shift results in a small rotation of the plane of linear polarization. We denote this rotation by $\Delta\theta$.

Simultaneously, the difference in absorption coefficients, $\alpha_+ \neq \alpha_-$, means one component is absorbed more strongly than the other. When they recombine, they no longer have equal amplitudes. This imbalance transforms the originally linear polarization into a slightly elliptical polarization.

A useful analogy here is the Faraday effect, where an external magnetic field aligns the atomic angular momenta, J, causing a polarization rotation. The beautiful thing about polarization spectroscopy is that we achieve this alignment *optically* with the pump beam. No external magnetic field is required.

Page 15:

This pair of diagrams provides an excellent visual summary of the transition from an isotropic to an anisotropic medium.

On the left, in panel (a), we see a depiction of an Isotropic Gas. We have a collection of molecules, and the red arrows represent the orientation of their individual angular momentum vectors, J J. As you can see, these arrows point in random directions. There is no preferred orientation. Macroscopically, this means the refractive index for σ + σ^+ light, n + n_+, is exactly equal to the refractive index for σ − σ^- light, n − n_-. The medium is optically uniform.

Now, look at panel (b), the Optically Pumped Anisotropic Gas. A strong, circularly polarized pump beam, represented by the thick green arrow, passes through the gas. Due to the selection rules we just discussed, this pump beam preferentially depletes certain M M-sublevels. The result is that the angular momentum vectors are no longer randomly oriented; they now have a preferred orientation, aligned with respect to the pump beam's propagation axis.

The lower part of panel (b) shows the consequence for our probe beam. The probe beam, composed of its σ + σ^+ (red wave) and σ − σ^- (blue wave) components, enters this anisotropic medium. Because n + n_+ is not equal to n − n_-, the two components travel at different speeds. The diagram beautifully illustrates that one wave gets phase-shifted relative to the other. When these two components recombine, the resulting polarization is rotated. This is the essence of birefringence, and it's the physical origin of our signal. The result is clear: n + n_+ is not equal to n − n_-.

Page 16:

Up to now, we've focused on the polarization aspect. But this is a Doppler-free technique. So, let's re-introduce the velocity of the atoms and see how the Doppler effect plays a crucial role.

As we've established, the pump and probe beams are counter-propagating. This means their wave vectors are equal and opposite: the k k-vector for the pump is equal to minus the k k-vector for the probe.

Now, consider an atom moving with a velocity component v z v_z along the laser beam axis. Due to the Doppler effect, this atom does not see the laser's lab-frame frequency, ω ω. It sees a shifted frequency. For the pump beam, this would be ω − k v z $\omega - k\,v_z$. For the probe, it would be ω + k v z $\omega + k\,v_z$. The atom will only interact resonantly with the light if this Doppler-shifted frequency matches its natural transition frequency, ω 0 ω_0.

So, if we set our laser to a frequency ω ω, which is slightly detuned from the line center ω 0 ω_0, the resonance condition will only be met for molecules with a specific axial velocity v z v_z. Rearranging the Doppler shift formula gives us, to first order, that the selected velocity is

v z ≈ ω 0 − ω k .

$$v_z \approx \frac{\omega_0 - \omega}{k}.$$

This means that the pump beam doesn't interact with all the atoms. At a given frequency, it selectively interacts with, and "burns a hole" in the population of, a very specific velocity class.

Page 17:

Let's be precise about the velocity classes for each beam.

The pump beam propagates, let's say, in the plus-z direction. The resonance condition for an atom with velocity v z v_z is that the atomic resonance frequency, ω 0 ω_0, must equal the Doppler-shifted laser frequency, ω − k v z $\omega - k\,v_z$. Solving for v z v_z, we find that the pump interacts with molecules having an axial velocity of v z = (ω − ω 0) / k $v_z = (\omega - \omega_0)/k$. The slide has a minus sign, so let's be consistent and define detuning as ω 0 − ω $\omega_0 - \omega$. So, the pump interacts with v z v_z equals *plus* (ω 0 − ω) / k $(\omega_0 - \omega)/k$. The small γ / k γ/k term represents the small range of velocities within the natural linewidth.

Now, what about the probe? It's counter-propagating, so its wave vector is − k $-k$. The resonance condition for the probe is ω 0 = ω − (− k) v z $\omega_0 = \omega - (-k)\,v_z$, which is ω + k v z $\omega + k\,v_z$. Solving for v z v_z, we find the probe interacts with molecules having v z v_z equals *minus* (ω 0 − ω) / k $(\omega_0 - \omega)/k$.

So, you see, when the laser is detuned from resonance (when ω ω is not equal to ω 0 ω_0), the pump and probe beams interact with two completely different, distinct velocity classes, symmetric about v z = 0 $v_z = 0$. The probe beam never encounters the atoms that have been polarized by the pump.

The crucial exception is when we tune the laser frequency, ω ω, to be exactly on resonance with the atomic transition, so that ω ω is approximately equal to ω 0 ω_0.

In this special case, the detuning is zero. The pump interacts with the v z v_z approximately equal to zero velocity class, and the probe *also* interacts with the v z v_z approximately equal to zero velocity class.

This is the key to the Doppler-free nature of the signal. It is only when both beams interact with the same molecules—the stationary or near-stationary ones—that the probe can "see" the anisotropy created by the pump. It is only then that the medium becomes birefringent for the probe, and we get a signal.

Page 18:

The direct and powerful consequence of this condition is that the technique produces a sharp, Doppler-free peak exactly at the line center, where ω = ω 0 $\omega = \omega_0$. When we scan the laser frequency across the entire Doppler-broadened profile, we will see a signal *only* at the precise moment we pass through the true, natural resonance frequency of the stationary atoms.

This allows us to measure transition frequencies with extremely high precision, free from the limitations of thermal motion.

Page 19:

This slide offers a fantastic graphical representation of the velocity selection mechanism. Let’s walk through it.

In both panels, the blue curve represents the Maxwell–Boltzmann distribution of axial velocities, v z v_z, in our gas sample. The peak is at v z = 0 $v_z = 0$, and the width represents the thermal spread of velocities that leads to Doppler broadening.

First, let’s examine the left panel, the “Off-Resonance Condition”, where the laser frequency, ω ω, is not equal to the atomic resonance, ω 0 ω_0. The laser detuning, (ω 0 − ω) / k $(\omega_0 - \omega)/k$, is non-zero. As we derived, the pump beam interacts with a specific velocity class, shown as the red slice, at a positive v z v_z. The counter-propagating probe beam interacts with a different velocity class, shown as the blue slice, at a negative v z v_z. They are interacting with completely separate populations of atoms. The atoms that the probe sees have not been affected by the pump. Therefore, no polarization signal is generated. All we would measure is the standard, broad Doppler-broadened absorption.

Now, turn your attention to the right panel, the “On-Resonance Condition”. Here, we have tuned the laser so that ω ω is approximately equal to ω 0 ω_0. The detuning is now essentially zero.

Both the pump and the probe beams interact with the *same* class of molecules: those with near-zero axial velocity, right at the center of the distribution. This region of overlap is shown in purple.

In this case, the intense pump beam aligns the molecules in this v z ≈ 0 $v_z \approx 0$ group, creating birefringence and dichroism. The probe beam then passes through this aligned group, its polarization gets rotated, and it generates a sharp, Doppler-free signal.

The slider at the bottom conceptually represents tuning the laser frequency, showing how the two selected velocity packets move symmetrically until they merge at the center.

Page 20:

Having established the physical principles, let's now build a more rigorous mathematical description of the signal. This will be a mathematical prelude to our final signal expression.

We begin by decomposing the probe field. Let's assume for concreteness that our initial probe beam, after passing through the first polarizer P 1 P_1, is linearly polarized along the x x-axis. We can write its electric field as a plane wave:

E → (z , t) = E 0 x ^ e i (ω t − k z) .

$$\vec{E}(z,t) = E_0\,\hat{x}\,e^{i(\omega t - kz)}.$$

Now, here is the essential mathematical step we discussed conceptually earlier. We will express this single linearly-polarized field as the sum of two equal-amplitude, counter-rotating circular components. We write:

E → = E → + + E → − .

$$\vec{E} = \vec{E}_+ + \vec{E}_-.$$

Where E → + $\vec{E}_+$ represents the right-hand circularly polarized component, and E → − $\vec{E}_-$ represents the left-hand circularly polarized component.

Page 21:

Here we see the explicit expressions for these circular components.

E-plus is equal to E 0 2 $\frac{E_0}{2}$, times the complex vector (x ^ + i y ^) $(\hat{x} + i\hat{y})$, all multiplied by the same plane wave factor, e i (ω t − k z) $e^{i(\omega t - kz)}$. The vector part, (x ^ + i y ^) $(\hat{x} + i\hat{y})$, describes a vector that rotates in the x-y plane and corresponds to right-hand circular or σ + σ^+ polarization.

Similarly, E-minus is equal to E 0 2 $\frac{E_0}{2}$, times the complex vector (x ^ − i y ^) $(\hat{x} - i\hat{y})$, times the same plane wave factor. The (x ^ − i y ^) $(\hat{x} - i\hat{y})$ term corresponds to left-hand circular or σ − σ^- polarization.

You can easily verify that if you add E + E_+ and E − E_-, the i y ^ $i\hat{y}$ terms cancel, and you are left with the original x-polarized field.

The profound advantage of this decomposition is that we can now analyze the interaction of our probe with the anisotropic medium by considering the two circular modes, σ + σ^+ and σ − σ^-, completely independently. The medium, having been prepared by the σ + σ^+ pump beam, has different properties—a different refractive index and a different absorption coefficient—for each of these two components.

So, all of our further calculations will be performed for each component separately. We'll calculate how E + E_+ propagates and how E − E_- propagates. Then, at the end, we will recombine them to find the final state of the field.

Page 22:

Let's now consider the propagation of our probe beam's components through the anisotropic sample. We will follow the logic of Slide 7.

Let's say the probe travels through a length L of the pumped medium. We'll focus first on the σ_+ component of the probe. As it propagates, two things happen: its phase evolves and its amplitude is attenuated. The refractive index for this component is n_+, which corresponds to a wave vector k_+. The absorption coefficient is α_+.

So, the field of the σ_+ component after traveling a distance L, which we write as $E_+(L)$, is given by the following expression:

E + (L) = E 0 2 (x ^ + i y ^) exp (i ω t − i k + L − α + L 2) .

$$E_+(L) = \frac{E_0}{2}(\hat{x} + i\hat{y})\exp(i\omega t - ik_+L - \alpha_+L/2).$$

$E_+(L)$ is equal to its initial amplitude, $\frac{E_0}{2}$ times $(\hat{x} + i\hat{y})$, multiplied by an exponential term. Inside the exponential, we have $\exp$ of the quantity: $i\omega t$ which is just the time oscillation, minus ik_+L, which is the phase accumulated due to propagation, minus $\alpha_+L/2$, which is the amplitude attenuation. Note that the absorption coefficient α appears in the field expression as $\alpha/2$.

An exactly analogous expression can be written for the E_- component, $E_-(L)$, but it will involve the wave vector k_- and the absorption coefficient α_-.

Now, we will define some terms to simplify our upcoming algebra.

Page 23:

The crucial physical quantities are the *differences* in how the medium treats the two circular polarizations. We define these on this slide.

First, the difference in refractive indices, which causes birefringence:

Δ n = n + − n − .

$$\Delta n = n_+ - n_-.$$

Second, the difference in absorption coefficients, which causes circular dichroism:

Δ α = α + − α − .

$$\Delta \alpha = \alpha_+ - \alpha_-.$$

As the probe's two circular components propagate through the length L L of the sample, they accumulate a relative phase difference. This resulting phase difference, Δ ϕ $\Delta\phi$, is given by:

Δ ϕ = (k + − k −) L .

$$\Delta \phi = (k_+ - k_-)\, L.$$

Since the wavevector k k is related to the refractive index n n by k = ω n / c $k = \omega n / c$, this becomes:

Δ ϕ = ω L c Δ n .

$$\Delta \phi = \frac{\omega L}{c}\, \Delta n.$$

This phase difference is what causes the polarization plane to rotate.

Simultaneously, the two components accumulate a resulting amplitude difference due to differential absorption. The difference in field amplitude, Δ E ΔE, after the sample is:

Δ E = E naught 2 (e − α + L / 2 − e − α − L / 2) .

$$\Delta E = \frac{E_{\text{naught}}}{2}\left(e^{-\alpha_{+}L/2} - e^{-\alpha_{-}L/2}\right).$$

This amplitude difference is what makes the polarization elliptical.

It is clear from this that both Δ ϕ $\Delta\phi$ and Δ E ΔE, which depend on Δ n Δn and Δ α $\Delta\alpha$ respectively, encode the information about the interaction. This is our spectroscopy signal.

Page 24:

Now we must address an important, and often unavoidable, experimental reality: the effect of the windows of the sample cell. This is not just a trivial detail; it can significantly impact the signal.

The cell windows, typically made of glass or quartz and having a thickness we'll call ‘d’, can themselves contribute extra, or “parasitic,” birefringence and absorption. This usually arises from mechanical stress induced during manufacturing or from the pressure difference when the cell is evacuated.

To handle this, we introduce a complex refractive index to describe the properties of the windows. We'll denote it by a tilde over the n. For the two circular polarizations, we have:

n ~ w ± = b r ± + i b i ± .

$$\tilde{n}_{w\pm} = b_{r\pm} + i\, b_{i\pm}.$$

Here, b r b_r is the real part of the window’s refractive index, and b i b_i is the imaginary part.

The imaginary part, b i b_i, is directly related to absorption. We can relate it to an absorption coefficient for the windows, as we'll see next.

Page 25:

The relationship between the imaginary part of the refractive index of the window, b i b_{i}, and the window's absorption coefficient, α w α_{w}, is given by:

α w ± = 2 ω c b i ± .

$$\alpha_{\mathrm{w}}^{\pm} = \frac{2\omega}{c}\, b_{\mathrm{i}}^{\pm}.$$

We can then define a total absorption for the windows, which we'll call ' a w a_{w}', that accounts for the fact that the probe passes through two windows (at the entrance and exit of the cell). If each window has thickness ' d d', the total path length in the glass is 2 d $2\,d$. So:

a w ± = 2 d α w ± .

$$a_{\mathrm{w}}^{\pm} = 2\, d\, \alpha_{\mathrm{w}}^{\pm}.$$

Just as we did for the gas sample, we are interested in the *differences* for the two circular polarizations. We define:

Δ b r = b r + − b r − .

$$\Delta b_{\mathrm{r}} = b_{\mathrm{r}}^{+} - b_{\mathrm{r}}^{-}.$$

This is the parasitic birefringence of the windows.

Δ a w = a w + − a w − .

$$\Delta a_{\mathrm{w}} = a_{\mathrm{w}}^{+} - a_{\mathrm{w}}^{-}.$$

This is the parasitic dichroism of the windows.

The crucial takeaway here is that these parasitic terms can distort, or worse, even mimic a real signal from the atoms. They introduce their own rotation and ellipticity. Therefore, for a high-quality experiment, the windows must be made from high-quality, strain-free material, and they must be carefully characterized or their effects compensated.

Page 26:

This diagram provides a clear illustration of how parasitic birefringence arises in the vapor cell windows.

On the left, we see the vapor or gas sample, and attached to it is the cell window, a cylinder of glass with thickness d d. Due to mechanical stress from how the window is manufactured or mounted, the glass is often not perfectly isotropic. It can develop privileged axes, known as the "fast axis" and the "slow axis," indicated by the red arrows. This means the refractive index is different for light polarized along these two directions.

As the text box below explains, this stress-induced anisotropy means that when the probe beam passes through, its two circular polarization components, σ + σ^{+} and σ − σ^{-}, experience different complex refractive indices. This results in a parasitic phase difference, related to Δ b Δb that we just defined, and a parasitic differential absorption, related to Δ a Δa.

This unwanted effect from the windows will add to the real signal from the gas. It's a source of systematic error that must be carefully managed in any precision polarization spectroscopy experiment.

Page 27:

Slide 9: Recombination & Emergent Elliptically Polarised Field

Alright, let's bring our mathematical description to its climax. We've calculated how the E-plus and E-minus components of the probe field propagate through both the pumped gas and the cell windows. Now it's time to recombine them and see what the final emergent field looks like. This is Slide 9, "Recombination & Emergent Elliptically Polarised Field".

After passing through the cell of length L, the total field, E-out, is simply the vector sum of the two components:

E o u t = E + (L) + E − (L)

$$E_{\text{out}} = E_{+}(L) + E_{-}(L)$$

The algebra to combine the expressions from Slide 7, including the window effects, is a bit tedious, but the result is beautifully compact. The recombination yields a field that is, in general, elliptically polarized. It can be described by the following equation:

E o u t = E 0 2 e i ω t e − i Φ 0 (x ^ + i y ^ e − i δ)

$$E_{\text{out}} = \frac{E_0}{2}\, e^{\,i\omega t}\, e^{-\,i\Phi_0}\left(\hat{x} + i\,\hat{y}\,e^{-\,i\delta}\right)$$

Let's look at this. Φ 0 Φ_0 (capital Phi sub zero) is a common phase factor that affects the whole field. The interesting part is the vector term. It is no longer x ^ + i y ^ $\hat{x} + i\hat{y}$, which would be circular, or a simple real vector, which would be linear. The presence of the complex phase factor e − i δ $e^{-i\delta}$ on the y-component is what describes the new elliptical polarization state. The complex phase δ δ (lowercase delta) contains all the physics of the interaction.

Page 28:

So what are these phase terms, Phi-naught and delta? Let's define them.

The common phase, capital Phi-naught, is given by:

Φ 0 = ω c n L + ω c b r d − i (α L 2 + a w 2) .

$$\Phi_0 = \frac{\omega}{c}\, nL \;+\; \frac{\omega}{c}\, b_r d \;-\; i\left(\frac{\alpha L}{2} + \frac{a_w}{2}\right).$$

This term represents the average phase shift and average absorption from both the gas (terms with n n and α α) and the windows (terms with b r b_r and a w a_w). This is generally not what we're interested in.

The crucial term is the complex phase difference, lowercase delta:

δ = ω L c Δ n + ω d c Δ b r − i (L Δ α 2 + Δ a w 2) .

$$\delta = \frac{\omega L}{c}\, \Delta n \;+\; \frac{\omega d}{c}\, \Delta b_r \;-\; i\left(\frac{L\,\Delta\alpha}{2} + \frac{\Delta a_w}{2}\right).$$

This term contains *all* the differential effects—the differences in refractive index and absorption for both the gas sample and the windows. This is our signal.

Let's break down δ δ.

The real part of δ δ, which comes from Δ n Δn and Δ b r Δb_r, is directly related to the rotation of the polarization axis. The rotation angle is approximately the real part of δ δ, divided by 2.

The imaginary part of δ δ, which comes from Δ α $\Delta\alpha$ and Δ a w Δa_w, is responsible for the differential attenuation between the σ + σ^+ and σ − σ^- components. This is what causes the probe to become elliptical.

So, by measuring the final polarization state, which is fully characterized by this complex number δ δ, we can measure the underlying physics of the light-matter interaction.

Page 29:

Slide 10:

We're now at the final step of the optical path: the analyzer. This is Slide 10. After the probe field emerges from the cell in its new, elliptically polarized state, E o u t E_{out}, it must pass through the second polarizer, P 2 P_2, before reaching the detector.

Ideally, P 2 P_2 is crossed with the initial polarizer P 1 P_1. If P 1 P_1 defines the x-axis, then P 2 P_2 would define the y-axis. However, for reasons that will become clear, it's often useful to deliberately "uncross" the analyzer by a small angle. Let's call this small uncrossing angle θ θ (theta). So, the transmission axis of P_2 is tilted by a small angle θ away from the y-axis.

The detector only measures the component of the E_{out} field that is projected onto the transmission axis of P_2. The transmitted field amplitude, which we'll call E_{t}, is given by simple vector projection. Using the small angle approximation where $\sin(\theta)$ is about θ and $\cos(\theta)$ is about 1, and given the analyzer is mostly along y, the transmitted field is:

E t = E x sin (θ) + E y cos (θ) .

$$E_{\text{t}} = E_{\text{x}}\sin(\theta) + E_{\text{y}}\cos(\theta).$$

Where E_{x} and E_{y} are the x and y components of our E_{out} field from the previous slide. This projection is what we will use to calculate our final signal.

Page 30:

To arrive at a clear and interpretable expression for our signal, we will now make the small-signal approximation. This is a very reasonable assumption in most experimental conditions, where the changes induced by the pump beam are indeed small perturbations.

We assume that all the differential quantities are much less than 1. Specifically: - The differential absorption in the gas, $L\Delta\alpha$, is much, much less than 1. - The differential phase shift in the gas, $\frac{\omega L\Delta n}{c}$, is much, much less than 1. - And similarly, the window birefringence Δb_{r} and dichroism Δa_{w} are also much, much less than 1.

Under this assumption, the complex phase δ we defined is a small quantity. This allows us to perform a Taylor series expansion of the exponential $\exp(-i\delta)$ that appeared in our expression for the output field. We expand it as: $\exp(-i\delta) \approx 1 - i\delta$.

We will keep terms only through the first order in all of these small quantities. This linearizes the problem and will give us a very clean final result for the transmitted field.

Page 31:

Applying the small-signal approximation to our expression for the transmitted field E_t gives us this wonderfully simple and insightful result:

E_t is approximately equal to E_0 times a complex phase factor $e^{i\omega t - i\Phi_0}$, all multiplied by the simple sum $(\theta + \delta)$.

Let's pause and appreciate this equation. It tells us that the amplitude of the light reaching the detector is proportional to the sum of two independent terms, our two "control knobs."

1. The first term is θ, the deliberate, mechanical uncrossing angle of the analyzer. This is a static, controllable parameter that we can set in the lab.
2. The second term is δ, the complex phase which contains all the interesting, pump-induced physics: the rotation and dichroism from our atomic sample, Δn and $\Delta \alpha$, as well as the parasitic window effects.

The final signal will arise from the interplay and interference of these two terms.

Page 32:

Now we can finally calculate the signal that our detector measures. Detectors like photodiodes are square-law detectors; they measure intensity, which is proportional to the electric field amplitude squared.

Slide 11: the general expression for the Detector Intensity

This brings us to Slide 11, the general expression for the Detector Intensity.

The intensity, I_{T}, is given by a constant times the permittivity of free space, ϵ_0, times the modulus squared of the transmitted field amplitude, E_{t}. Using our result from the last slide, this means the intensity is proportional to the modulus squared of the quantity $(\theta + \delta)$. So, I_{T} is proportional to $|\theta + \delta|^2$.

But we have one more piece of experimental reality to include. Even the best polarizers are not perfect. When they are ideally crossed, they don't block 100% of the light. There is always a small "residual transmission" or "leakage." We characterize this by the extinction ratio, ξ (the Greek letter xi), which is defined as the ratio of the residual intensity transmitted through crossed polarizers, I_{res}, to the incident intensity, I_0. Typically, for good polarizers, ξ is a very small number, on the order of 10^{-6} to 10^{-8}. This ξ term will act as a fundamental background floor in our measurement.

Finally, to simplify the final expression, it's convenient to define a "shifted angle," θ ′ θ' (theta-prime), which absorbs the static window birefringence term. We define:

θ ′ = θ + ω 2 c Δ b r .

$$\theta' = \theta + \frac{\omega}{2\,c}\,\Delta b_{\mathrm{r}}.$$

This combines the mechanical uncrossing angle with the static window rotation into a single parameter.

Page 33:

Putting all of these pieces together—the squared amplitude, the finite extinction ratio ξ ξ, and the parasitic window effects—we arrive at the final expression for the total detected signal, I T (ω) $I_{\mathrm{T}}(\omega)$, as a function of the laser frequency omega. The equation looks a bit complex, but we can break it down:

I T (ω) = I 0 e − α L − a w [ξ + θ ′ 2 + (1 2 Δ a w) 2 + 1 2 θ ′ L Δ α (ω) + ω 2 c θ ′ L Δ n (ω) + …]

$$I_{\mathrm{T}}(\omega) = I_0 e^{-\alpha L - a_{\mathrm{w}}}\left[\xi + \theta'^2 + \left(\frac{1}{2}\Delta a_{\mathrm{w}}\right)^2 + \frac{1}{2}\theta' L \Delta\alpha(\omega) + \frac{\omega}{2\,c}\theta' L \Delta n(\omega) + \cdots\right]$$

Let's dissect this.

- The I 0 e − α L − a w $I_0 e^{-\alpha L - a_{\mathrm{w}}}$ term out front is just the overall transmitted intensity, accounting for the average absorption in the gas and windows. - Inside the brackets, the first three terms (ξ ξ, θ ′ 2 θ'^2, and the Δ a w Δa_{w} squared term) are all independent of the laser frequency detuning. They form a constant DC background. - The last two terms are the ones we care about. They are the frequency-dependent contributions that constitute our spectroscopic signal.

Let's highlight these two contributions of interest:

1. A "Dispersion-type term," which is proportional to θ' times $\Delta n(\omega)$. But through the Kramers-Kronig relations, this is also related to $\Delta\alpha(\omega)$. This term's shape will be dispersive, or anti-symmetric. 2. A "Lorentzian-type term," which arises from the cross-term between the gas signal $\Delta\alpha(\omega)$ and the window dichroism Δa_{w}. Its shape will be Lorentzian, or symmetric.

The relative strength of these two terms can be controlled by our experimental parameters, θ' and Δa_{w}.

Page 34

To understand the shape of our signal, we need a model for the frequency dependence of the differential absorption, $\Delta\alpha(\omega)$. This is the subject of Slide 12.

The key physical insight is that the signal is generated *only* by those molecules that are simultaneously in resonance with both the pump and the probe beams. As we established, this only happens for the $v_{\mathrm{z}} \approx 0$ velocity class, when the laser is tuned near the line center ω_0.

Therefore, the lineshape of our signal is not the Doppler-broadened Gaussian profile. Instead, it's the natural, homogeneous lineshape of the transition, which is a Lorentzian. The width of this Lorentzian is the homogeneous half-width, γ s γ_s (gamma-sub-s). This width includes natural broadening, collisional broadening, and, importantly, power broadening from the strong pump beam.

To describe the lineshape mathematically, it's convenient to define a dimensionless detuning parameter, which we'll call x x.

x x is defined as the frequency detuning from line center, (ω 0 − ω) $(\omega_0 - \omega)$, normalized by the homogeneous half-width at half-maximum, which is γ s / 2 $\gamma_\mathrm{s}/2$.

x = ω 0 − ω γ s / 2

$$x = \frac{\omega_0 - \omega}{\gamma_\mathrm{s}/2}$$

When x = 0 $x = 0$, we are at the exact line center. x = 1 $x = 1$ means we are detuned by one half-width.

With this definition, the Lorentzian change in the absorption coefficient has a very simple form.

Page 35:

The frequency-dependent differential absorption coefficient, Δ α (ω) $\Delta\alpha(\omega)$, is given by a classic Lorentzian profile:

Δ α (ω) = Δ α 0 1 + x 2

$$\Delta\alpha(\omega) = \frac{\Delta\alpha_0}{1 + x^2}$$

Let's define the terms here: - $\Delta\alpha_0$ (Delta alpha naught) is the peak differential absorption that occurs at the line center, where $x = 0$. It represents the maximum change in absorption induced by the pump. - Its units are typically inverse centimeters, or wavenumbers. - And x is the dimensionless detuning we just defined.

This symmetric, bell-shaped Lorentzian function describes how the absorption part of our signal varies as we scan the laser frequency across the resonance.

Page 36:

Now we move to the other crucial component of our signal: the dispersion part, which arises from the differential refractive index, $\Delta n(\omega)$. This is the topic of Slide 13.

One of the most beautiful and fundamental principles in physics is causality, which leads to the Kramers–Kronig relations. These relations state that the real and imaginary parts of the linear response function of a system—in our case, the susceptibility, χ (chi)—are not independent. The real part of χ gives the refractive index, and the imaginary part gives the absorption. If you know one of them over all frequencies, you can, in principle, calculate the other.

We don't need to perform the full integral here. Using the standard dispersion integral for a single Lorentzian absorption profile, we can directly write down the corresponding change in refractive index, $\Delta n(\omega)$. The result is:

Δ n (ω) = c ω 0 Δ α 0 x 1 + x 2 .

$$\Delta n(\omega) = \frac{c}{\omega_0} \Delta\alpha_0 \frac{x}{1 + x^2}.$$

Let's examine the important features of this dispersion profile:

1. Notice the factor of x x in the numerator. This makes the function *odd* in x x. It's an antisymmetric profile. It's positive on one side of the resonance and negative on the other. 2. Because it's an odd function, it has a zero-crossing exactly at the line center, where x = 0 $x = 0$, which corresponds to ω = ω 0 $\omega = \omega_0$.

This zero-crossing is what makes the dispersion signal so incredibly useful for laser frequency locking, as it provides a perfect, unambiguous lock point.

Page 37:

We now have all the ingredients in hand. We have the general expression for the detector intensity from Slide 11, and we have the specific functional forms for the Lorentzian absorption Δ α (ω) $\Delta\alpha(\omega)$ and the dispersive refractive index Δ n (ω) $\Delta n(\omega)$ from the last two slides.

The task now, as outlined in Slide 14, is to substitute these results back into our intensity expression to get the complete, final signal shape for the case of a circularly polarized pump. This will give us a comprehensive formula that describes what we actually measure in the experiment as we scan the laser frequency.

Page 38:

Here it is, the complete expression for the signal, S c p (ω) $S^{cp}(\omega)$, where 'cp' stands for circularly polarized pump. It looks a bit formidable, but we will break it down term by term.

S^{cp}(\omega) = I_0 \cdot e^{-\alpha L - a_\text{w}} \left\{
\left[\xi + \theta'^2 + \tfrac{1}{4}\,\Delta a_\text{w}^{2\right]}
+ \tfrac{1}{2}\,\theta'\,\Delta\alpha_0 L \cdot \left[\frac{x}{1 + x^2}\right]
+ \left[\tfrac{1}{4}\,\Delta a_\text{w}\,\Delta\alpha_0 L + \left(\tfrac{\Delta\alpha_0 L}{4}\right)^2\right] \cdot \left[\frac{1}{1 + x^2}\right]
+ \tfrac{3}{4} \cdot \left[\frac{\Delta\alpha_0 x}{1 + x^2}\right]^2 \right\}

$$S^{cp}(\omega) = I_0 \cdot e^{-\alpha L - a_\text{w}} \left\{ \left[\xi + \theta'^2 + \tfrac{1}{4}\,\Delta a_\text{w}^{2\right]} + \tfrac{1}{2}\,\theta'\,\Delta\alpha_0 L \cdot \left[\frac{x}{1 + x^2}\right] + \left[\tfrac{1}{4}\,\Delta a_\text{w}\,\Delta\alpha_0 L + \left(\tfrac{\Delta\alpha_0 L}{4}\right)^2\right] \cdot \left[\frac{1}{1 + x^2}\right] + \tfrac{3}{4} \cdot \left[\frac{\Delta\alpha_0 x}{1 + x^2}\right]^2 \right\}$$

Let's analyze the behavior of each term, as suggested by the slide.

The first line, inside the curly braces, contains ξ ξ, θ ′ 2 θ'^2, and the Δ a w 2 Δa_w^2 term. None of these depend on the laser frequency detuning x x. This is our static, DC background signal.

Page 39:

Let's continue dissecting our signal equation.

The second line contains the term proportional to $\frac{x}{1+x^2}$. As we saw from the Kramers–Kronig relations, this is the pure **dispersion shape**. It is antisymmetric about the line center. This part of the signal is primarily controlled by the uncrossing angle, θ'.

The third line has terms proportional to $\frac{1}{1+x^2}$. This is the pure **Lorentzian shape**, which is symmetric about the line center. This part of the signal is primarily driven by the window dichroism, Δa_{w}, and also by a term quadratic in the sample absorption itself.

The fourth line is proportional to the square of the dispersion shape. This is a higher-order term. In the typical small-signal limit where $\Delta\alpha_0 L$ is much less than 1, this term is usually negligible and can be ignored.

So, our total signal is a sum of a constant background, a dispersion component, and a Lorentzian component.

Page 40:

In many applications, particularly for laser locking, the goal is to obtain a clean, pure dispersion signal. So, how can we achieve that? Slide 15 tells us how to set up the experiment to isolate the dispersion term.

Looking back at our signal equation, the dispersion term is multiplied by θ', and the main Lorentzian term is multiplied by Δa_{w}, the window dichroism. To isolate the dispersion term, we need to make its pre-factor large and the Lorentzian's pre-factor small.

Therefore, the condition is to set Δ a w Δa_{w} approximately to zero, and θ ′ θ' to be non-zero.

Practically, how do we do this?

1. To make Δ a w Δa_{w} near zero, we must minimize the stress in the cell windows. This can be done by using high-quality, strain-free glass, and by carefully mounting the windows. Sometimes, experimenters will even gently squeeze the cell with a clamp to actively compensate for the birefringence caused by the pressure difference between the inside and outside of the cell.

2. To make θ ′ θ' non-zero, we simply need to deliberately uncross the polarizers P1 and P2 by a small angle, θ θ. The optimal angle, as we will see when we discuss signal-to-noise, is typically chosen to be on the order of the square root of the extinction ratio, ξ ξ.

Page 41

What is the outcome of setting up the experiment this way?

The outcome is that the term proportional to θ ′ Δ α 0 $\theta' \Delta\alpha_0$ becomes the dominant frequency-dependent part of the signal. This yields a nearly pure dispersion peak, with its characteristic anti-symmetric shape and zero-crossing at the line center.

And why is this so useful? As I've mentioned, it's perfect for laser-frequency locking. The steep, linear slope of the dispersion signal as it passes through zero acts as an ideal error signal. If the laser frequency drifts slightly off resonance, a positive or negative voltage is generated. This voltage can be fed back into the laser's control electronics to push the frequency back to the exact line center. This provides a robust and stable lock without the need for additional frequency modulation, a technique known as frequency dithering, which can introduce its own complications.

Page 42:

Slide 16

Now let's consider a variation of the experiment. What happens if we use a linearly polarized pump beam instead of a circular one? This is the topic of Slide 16.

When the pump beam is linearly polarized (let's say oriented at 45 degrees to the probe's x-axis for maximum effect), the angular momentum selection rule changes. For a pump polarized along the quantization axis, the selection rule is Δ M = 0 $\Delta M = 0$. This creates a different kind of anisotropy in the medium, known as alignment rather than orientation.

We can perform an analogous derivation, starting with this new selection rule. I won't go through all the steps, but the result is a modified signal expression. The signal for a linearly polarized pump, S L P (ω) $S^{LP}(\omega)$, is given by:

S L P (ω) = I 0 e − α L − a w { [ξ + 1 4 θ 2 Δ a w 2 + (ω 2 c Δ b r) 2] + Δ b r 4 ω c Δ α 0 L [x 1 + x 2] + [− 1 4 θ Δ a w Δ α 0 L + (Δ α 0 L 4) 2] [1 1 + x 2] } .

$$
\begin{aligned}
S^{LP}(\omega) \quad &= I_0\, e^{-\alpha L - a_w}\{[\xi + 1/4\,\theta^2\,\Delta a_w^2 + (\omega/2\,c\ \Delta b_r)^2] \\
&\quad + \frac{\Delta b_r}{4}\frac{\omega}{c}\,\Delta\alpha_0 L\,[x/1 + x^2] \\
&\quad + [-\ 1/4\,\theta\,\Delta a_w\,\Delta\alpha_0 L + (\Delta\alpha_0 L/4)^2]\,[1/1 + x^2]\}.
\end{aligned}
$$

Page 43:

Let's compare this new signal expression for a linear pump with the one we derived for a circular pump.

The most striking feature is an interchange of roles between the experimental parameters.

- In the circular pump case, the dispersion term was proportional to the analyzer uncrossing angle θ θ. - Now, in this linear pump case, look at the dispersion term (the one with x / (1 + x 2) $x/(1+x^2)$). Its coefficient is proportional to Δ b r Δb_r, the *window birefringence*. - Conversely, the Lorentzian term (the one with 1 / (1 + x 2) $1/(1+x^2)$) is now proportional to the analyzer angle θ θ and the window dichroism Δ a w Δa_w.

This is a fascinating and important result. The choice of pump polarization fundamentally changes how different experimental imperfections and parameters contribute to the final lineshape. Understanding this is key to correctly interpreting your data and designing your experiment.

Page 44:

So far, we have implicitly assumed a simple two-level system, where the pump and probe lasers are one and the same, driving a single transition between a lower level (which might have degenerate sublevels) and an upper level. However, the technique is more versatile than that. This brings us to pump-probe level schemes beyond two levels.

By using two independent, tunable lasers—one for the pump and one for the probe—we gain enormous flexibility. We can explore a variety of level schemes.

1. The **Two-level** scheme is the standard one we've been discussing. The pump and probe have the same frequency and interact with a common lower and upper level, including their degenerate M-sublevels.

2. A **V-type** scheme involves one common lower level and two different upper levels. The pump laser drives the transition to one upper level, while the probe laser drives the transition from the *same* lower level to a *different* upper level.

3. A **Lambda-type** scheme, denoted by the Greek letter Lambda, involves two different lower levels and one common upper level. Here, the pump might excite atoms from one lower level to the common upper level, and the probe would measure a transition from a *different* lower level to that *same* upper level.

Page 45:

The key takeaway from these different schemes is this: Polarization spectroscopy can be used to probe any transition that shares *either the lower or the upper state* with the pumped transition.

Why does this work? Let's think about it. In a V-type scheme, the pump creates an anisotropic population distribution in the shared *lower* state. The probe, being resonant with a transition out of that same lower state, will "see" this anisotropy and generate a signal.

In a Lambda-type scheme, the pump depletes the population of a lower state by moving atoms to the shared *upper* state. This can create an oriented population in the upper state. If the probe is resonant with a transition into that same upper state, it will also experience a modified, anisotropic environment and generate a signal.

This flexibility makes polarization spectroscopy a powerful tool for exploring connections and interactions between different energy levels in complex atoms and molecules.

Page 46:

This slide provides simple energy level diagrams to visualize the schemes we just discussed.

On the left, we have the **Two-Level System**. A single lower level and a single upper level are shown. Both the pump (pink arrow) and the probe (blue arrow) are resonant with this same transition.

In the center is the **V-Type System**. It looks like the letter 'V'. There is one common lower level. The pump excites the system to one upper level, while the probe beam monitors a transition to a second, different upper level.

On the right is the **Lambda-Type System**, which looks like an inverted 'V'. Here we have two distinct lower levels and one common upper level. The pump excites the system from one of the lower levels, and the probe examines a transition from the other lower level up to that same common upper level. This scheme is particularly important in fields like coherent population trapping and electromagnetically induced transparency.

Page 47:

Slide 18: Saturation Parameter

Let's now dig deeper into the quantitative aspects of the interaction. Slide 18 introduces the Saturation Parameter, S S, which quantifies the strength of the pump beam's interaction with the atoms.

The dimensionless saturation parameter, S S, is defined as:

S = 8 σ J J 1 M I 2 γ s R ∗ ℏ ω .

$$S = \frac{8\,\sigma_{JJ_1M}\,I_2}{\gamma_{\mathrm{s}}\,R^{*}\,\hbar\,\omega}.$$

Let's break down these terms: - σ J J 1 M σ_{JJ_1M} is the absorption cross-section for the specific pump transition out of the sublevel | J , M ⟩ $|J, M\rangle$. It has units of area, like centimeters squared. - I 2 I_2 is the intensity of the pump beam, in units like Watts per square centimeter. - γ s γ_{s} is the saturated homogeneous linewidth, which we've encountered before. It has units of inverse seconds. - R ∗ R^{*} is the population-relaxation rate out of the level, accounting for all decay channels. - ℏ ω $\hbar\omega$ is, of course, the energy of a single photon.

When S S is much less than 1, we are in the weak-pumping regime. When S S is on the order of 1 or greater, the pump is strong enough to significantly deplete the ground state population, and we say the transition is saturated.

Page 48:

The saturation parameter S S directly tells us how the population in a given sub-level is affected by the pump beam. After the pump has been applied, the new, saturated population, which we'll call N M S N_M^S, is related to the initial population, N M 0 N_{M0}, by this simple formula:

N M S = N M 0 1 + S .

$$N_M^S = \frac{N_{M0}}{1+S}.$$

So, if S = 1 $S=1$, the population is cut in half. If S S is very large, the population approaches zero.

This allows us to write a more fundamental expression for the peak differential absorption, Δ α 0 $\Delta\alpha_0$, that we introduced earlier. It can be shown that Δ α 0 $\Delta\alpha_0$ is equal to the unsaturated absorption coefficient, α 0 α_0, times the on-resonance saturation parameter, S 0 S_0, times a purely geometric factor, which we denote as Δ C J J 1 ∗ $\Delta C_{JJ_1}^*$.

Δ α 0 = α 0 S 0 Δ C J J 1 ∗ .

$$\Delta\alpha_0 = \alpha_0\, S_0\, \Delta C_{JJ_1}^*.$$

This expression is very powerful. It separates the physics into three parts: α 0 α_0, which depends on the total number of atoms; S 0 S_0, which depends on the laser power and intrinsic atomic properties; and Δ C ΔC, which depends only on the angular momentum quantum numbers of the transition.

Page 49:

To be clear about the terms in our new expression for Δ α 0 $\Delta\alpha_0$:

- S 0 S_0 is the saturation parameter evaluated at the line center, where the interaction is strongest. - Δ C J J 1 ∗ $\Delta C^*_{JJ_1}$ (Delta C star sub J J one) is a purely geometric factor that encapsulates all the angular momentum algebra. It depends on the J J values of the levels, the type of transition (P P, Q Q, or R R branch), and the polarization of the pump beam. We will calculate this factor next. It's what determines the relative strengths of the signals for different transitions.

Page 50:

This brings us to the core of the geometric factor: the Clebsch–Gordan weighted cross-sections, as described on Slide 19.

The absorption cross-section for a probe photon with a specific circular polarization (σ + σ^+ or σ − σ^-) on a transition from a specific M M sub-level is not uniform. We can write the differential absorption per sub-level, σ J J 1 M ± $\sigma^{\pm}_{JJ_1M}$, as:

σ J J 1 M ± = σ J J 1 × C (J , J 1 , M , M ± 1) .

$$\sigma^{\pm}_{JJ_1M} = \sigma_{JJ_1} \times C(J, J_1, M, M \pm 1).$$

Let's break this down:

- σ J J 1 σ_{JJ_1} is the orientation-averaged cross-section for the transition. This is the value you would measure if the sample were isotropic.

- C C is the square of the appropriate Clebsch–Gordan coefficient for the specific transition from sub-level M M to sub-level M ± 1 $M \pm 1$. Clebsch–Gordan coefficients are the mathematical tools from quantum mechanics that tell us how to add angular momenta. They determine the relative probabilities of transitions between different M M sub-levels. They are, in essence, the "rules" of angular momentum conservation made quantitative.

Page 51

To find the total macroscopic differential absorption, Δ α $\Delta\alpha$, we need to sum the contributions from all the M sub-levels, taking into account the population changes induced by the pump.

The total differential absorption for a transition from a manifold J to a manifold J1 is given by the sum over all M of: The population in state M, N M N_M, times the difference in cross-sections, σ J J 1 M + − σ J J 1 M − $\sigma^{+}_{JJ_1M} - \sigma^{-}_{JJ_1M}$.

This involves a lengthy summation of Clebsch–Gordan coefficients. Fortunately, this algebra has been worked out, and the result can be encapsulated in that single numerical factor, \(\Delta C^_*{J J_1}\), that we introduced. This factor essentially performs the weighted sum for us. The values of \(\Delta C*^\) have been tabulated and plotted for various cases, which we'll see on the next slide. This saves us from having to re-calculate these sums for every new experiment.

Page 52:

So, let's look at what this geometric factor, Δ C J J 1 $*$ $\Delta C^*_{JJ_1}$, depends on. This is outlined on Slide 20.

There are two separate cases to consider, based on the pump polarization, as this determines the selection rules.

1. **Linear pump**: This corresponds to the selection rule Δ M = 0 $\Delta M = 0$. 2. **Circular pump**: This corresponds to the selection rule Δ M = ± 1 $\Delta M = \pm 1$.

Within each of these cases, the value of Δ C $*$ ΔC^* also depends heavily on the type of rotational branch the transition belongs to. This is determined by the change in the J quantum number, Δ J ΔJ:

- **P-branch**: Δ J = − 1 $\Delta J = -1$. - **Q-branch**: Δ J = 0 $\Delta J = 0$. - **R-branch**: Δ J = + 1 $\Delta J = +1$.

So, for any given transition, we can look up or calculate the value of Δ C $*$ ΔC^* based on the pump polarization we choose and the branch type of the transition.

Page 53:

There is one more physical parameter that the geometric factor Δ C $*$ ΔC^* depends on, and that is the ratio of the population relaxation rates of the upper and lower levels. This is characterized by the parameter r r, defined as:

r = γ J − γ J 1 γ J + γ J 1

$$r = \frac{\gamma_{\mathrm{J}} - \gamma_{J_1}}{\gamma_{\mathrm{J}} + \gamma_{J_1}}$$

where γ J γ_{J} and γ J 1 γ_{J_1} are the relaxation rates of the two levels involved.

However, for our purposes, the most important information is the qualitative take-away, which is immensely powerful for practical spectroscopy.

1. A **linear pump** configuration strongly favors and enhances signals from **Q-branch** transitions. 2. A **circular pump** configuration strongly favors and enhances signals from **P-branch and R-branch** transitions.

This provides an incredibly powerful diagnostic tool. If you have a complex, congested spectrum with many overlapping lines and you don't know which transition is which, you can simply record the spectrum twice: once with a linear pump and once with a circular pump. The lines that are strong in the first spectrum are Q-lines. The lines that are strong in the second are P or R lines. This selectivity is one of the most celebrated features of polarization spectroscopy.

Page 54:

These graphs beautifully illustrate the selectivity we just discussed. They plot the calculated relative signal strength, which is proportional to our Δ C ∗ ΔC^* factor, as a function of the rotational quantum number J J.

Let's look at the top graph, for **Linear Pump Polarization**. The vertical axis is the relative signal, and the horizontal axis is J J. The blue line represents the signal for a Q-branch ($\Delta J = 0$). Notice that it is large and relatively constant for all but the very lowest J values. Now look at the red and green lines for the P- and R-branches. Their signals are much, much smaller. This plot visually confirms that a linear pump dramatically enhances Q-branch signals.

Now, let's examine the bottom graph, for **Circular Pump Polarization**. The situation is completely reversed. The red line (P-branch, $\Delta J = -1$) and the green line (R-branch, $\Delta J = +1$) show large signals that grow with increasing J. In contrast, the blue line for the Q-branch is much smaller in magnitude. This confirms that a circular pump is the right choice for observing P- and R-branch transitions. These plots are the theoretical basis for the powerful spectral assignment capability of the technique.

Page 55:

Let's look at a real-world application of this principle. Slide 21 shows an example from the spectrum of the Cesium dimer molecule, Cs_2 , recorded around a wavelength of 627.8 nanometers.

Two spectra were recorded of the exact same spectral region.

- The **upper trace** was taken using a **linear pump**. The result is exactly as the theory predicts: the Q-lines, which are transitions with $\Delta J = 0$, dominate the spectrum, while the P and R lines are very weak. - The **lower trace** was taken using a **circular pump**. Again, in perfect agreement with theory, the P and R lines are now strongly enhanced, while the Q-lines have all but disappeared, appearing only as small, residual dispersive features.

This is a stunning demonstration of the selectivity of polarization spectroscopy. It allows an experimenter to effectively "turn on" and "turn off" different types of transitions simply by rotating a half-wave plate to change the pump's polarization from linear to circular. This is an invaluable technique for assigning and understanding complex molecular spectra.

Page 56:

Here we see the actual data from the Cesium dimer example. The plot shows signal intensity in arbitrary units versus wavelength in nanometers.

The **Upper Trace**, for linear pump polarization, is shown in blue. You can see a very strong, sharp feature right in the middle, which is clearly labeled as the dominant Q-lines. To the sides, where the P and R lines would be, the signal is essentially flat and weak.

The **Lower Trace**, for circular pump polarization, is shown in red. The change is dramatic. The central Q-line signal is now suppressed, appearing as just a small residual wiggle. Instead, two massive signals have appeared on either side, labeled as "P/R Lines Enhanced." These correspond to the P- and R-branch transitions.

This figure is a perfect textbook illustration of how the choice of pump polarization acts as a powerful filter, allowing us to disentangle different components of a complex spectrum.

Page 57:

Having thoroughly explored the physics and applications, let's return to one of the technique's key advantages: the signal-to-noise ratio. Let's analyze the case for a dispersion signal, which is often the most useful.

The dominant source of noise in our measurement is typically fluctuations in the detector background. The background intensity, I b g I_{bg}, comes from the residual transmission of the polarizers, ξ ξ, and the deliberate uncrossing, θ θ. So, the background is:

I b g = I 0 · e − α L − a w · (ξ + θ 2)

$$I_{\mathrm{bg}} = I_0 \cdot e^{-\alpha L - a_{\mathrm{w}}} \cdot (\xi + \theta^2)$$

.

The signal we want to measure is the peak-to-peak dispersion amplitude, which we'll call Δ S max ΔS_{max}. From our previous formulas, this is given by:

Δ S max = I 0 · e − α L − a w · θ · Δ α 0 · L

$$\Delta S_{\mathrm{max}} = I_0 \cdot e^{-\alpha L - a_{\mathrm{w}}} \cdot \theta \cdot \Delta\alpha_0 \cdot L$$

.

To quantify the noise, we can define an intensity-noise coefficient, 'a', such that the noise on our incident laser beam is I 0 / a I_0/a. For example, if the laser has 1% intensity noise, then 'a' would be 100. This noise will propagate through to our background measurement.

Page 58:

Now we can write down the signal-to-noise ratio, S / N S/N. The signal is Δ S max $\Delta S_{\max}$ and the noise is the background I b g I_{bg} divided by our noise coefficient a a. Hence, after canceling common terms, we get:

S N = a θ Δ α 0 L ξ + θ 2 .

$$\frac{S}{N} = a\,\frac{\theta\,\Delta\alpha_0 L}{\xi + \theta^2}.$$

Now, we can optimize this! The angle θ θ is a parameter we can control. To find the maximum possible signal-to-noise ratio, we take the derivative of this expression with respect to θ θ and set it to zero. A quick calculation shows that the optimum occurs when θ 2 = ξ $\theta^2 = \xi$, or when the uncrossing angle θ θ is equal to the square root of the polarizer extinction ratio. This is a very important experimental rule of thumb.

Plugging this optimal θ θ back into the expression gives the maximum achievable signal-to-noise ratio:

(S / N) max = a Δ α 0 L 2 ξ .

$$(S/N)_{\max} = a\,\frac{\Delta\alpha_0 L}{2\sqrt{\xi}}.$$

Let's compare this to the S/N for saturation spectroscopy, which can be shown to be (S / N) sat = 1 2 a α 0 L $(S/N)_{\text{sat}} = 1/2\ a\,\alpha_0 L$. The ratio of these two gives us the gain factor of our technique.

Page 59:

After accounting for the geometric factors, the final expression for the gain in sensitivity of polarization spectroscopy over saturation spectroscopy is:

Gain = Δ C J J 1 ∗ 2 ξ .

$$\text{Gain} = \frac{\Delta C_{JJ_1}^*}{2\sqrt{\xi}}.$$

Let's look at this beautiful result. The gain depends on two things: 1. Δ C ∗ ΔC^*, the geometric factor, which is typically on the order of 1. 2. ξ ξ, the extinction ratio of our polarizers.

Since ξ ξ is a very small number, like 10 − 6 10^{-6} or 10 − 8 10^{-8}, the square root of ξ ξ is also small (10 − 3 10^{-3} or 10 − 4 10^{-4}). This means the gain factor can be enormous. If ξ ξ is 10 − 6 10^{-6}, the gain is on the order of 1 2 × 10 − 3 $1/2 \times 10^{-3}$, which is 500. This is where the huge sensitivity improvement comes from.

Page 60:

Slide 23

Let's briefly consider the signal-to-noise ratio for the other case: generating Lorentzian signals. This is Slide 23.

In this case, we keep the polarizers perfectly crossed, so our effective angle θ ′ θ' is zero, but we assume we have a finite window dichroism, Δ a w Δa_{w}. It's also possible to generate Lorentzian signals by carefully tuning this window birefringence, for example by squeezing the cell.

From our general signal formula, the maximum signal amplitude in this case, Δ S max ΔS_{max}, is:

Δ S max = I 0 · e − α L − a w · Δ α 0 L 4 · (Δ a w + 1 4 Δ α 0 L) .

$$\Delta S_{\text{max}} = I_0 \cdot e^{-\alpha L - a_{\text{w}}} \cdot \frac{\Delta \alpha_0 L}{4} \cdot (\Delta a_{\text{w}} + 1/4\ \Delta \alpha_0 L).$$

The background intensity, I b g I_{bg}, is now:

I b g = I 0 · e − α L − a w · (ξ + 1 4 Δ a w 2) .

$$I_{\text{bg}} = I_0 \cdot e^{-\alpha L - a_{\text{w}}} \cdot (\xi + 1/4\ \Delta a_{\text{w}}^{2).}$$

We can again form the ratio of signal to noise.

Page 61:

We can now optimize the signal-to-noise ratio for the Lorentzian case, this time by varying the window dichroism, Δ a w Δa_{w}. The optimal value for the window dichroism, Δ a w , opt $\Delta a_{w,\text{opt}}$, is found to be approximately 4 ξ / (Δ α 0 L) $4\xi/(\Delta \alpha_0 L)$, for the case where ξ ξ is much smaller than the absorption.

If we insert this optimal value back into the S / N $\mathrm{S/N}$ expression, we get a lengthy formula, but the most important feature is how it scales. The maximum S / N $\mathrm{S/N}$ is proportional to (Δ α 0 L) / ξ $(\Delta \alpha_0 L)/\xi$. Notice that this scales as one over ξ ξ, not one over the square root of ξ ξ as in the dispersion case. This suggests that even larger enhancements are possible.

The practical takeaway is that with good quality windows, where the intrinsic extinction ξ ξ is less than or equal to 10 − 8 10^{-8}, and with proper tuning of the window birefringence, it is possible to achieve sensitivity improvement factors of 1,000 to 10,000 over conventional saturation spectroscopy. This is truly remarkable.

Page 62:

Let's make this concrete with a worked example using some realistic numbers.

Slide 24:

Let's assume the following parameters for our experiment: 1. The geometric factor, Δ C ∗ ΔC^*, is 0.5. A typical value. 2. The polarizer extinction ratio, ξ ξ, is 10 − 6 10^{-6}. This corresponds to reasonably good, but not exceptionally expensive, polarizers. 3. The unsaturated line-center absorption, α 0 L $\alpha_0 L$, is 10 − 2 10^{-2}, or 1%. This is a fairly weak transition. 4. The on-resonance saturation parameter, S 0 S_0, is 0.1. We are weakly saturating the transition. 5. The laser intensity noise coefficient, a a, is 10 2 10^2, or 100, which corresponds to 1% intensity noise.

Page 63:

Now let's calculate the expected signal-to-noise ratios.

First, for our **dispersion signal** in polarization spectroscopy, we use the formula (S / N) p o l = a α 0 L S 0 Δ C ∗ 2 ξ $(S/N)_{pol} = \frac{a\,\alpha_0 L S_0 \Delta C^*}{2\sqrt{\xi}}$.

Plugging in the numbers... we have 100 × 10 − 2 × 0.1 × 0.5 2 10 − 6

$\frac{100\times10^{-2}\times0.1\times0.5}{2\sqrt{10^{-6}}}$, which simplifies to 0.05 2 × 10 − 3 = 25 $\frac{0.05}{2\times10^{-3}} = 25$. So, we expect a signal-to-noise ratio of 25. That's a very clean signal. The slide shows an alternate calculation using Δ α 0 $\Delta\alpha_0$, leading to the same result.

Now, for the standard **saturation spectroscopy signal**, the formula is (S / N) s a t = 1 2 a α 0 L S 0 $(S/N)_{sat} = \frac{1}{2}\,a\,\alpha_0 L S_0$. Plugging in the numbers: 0.5 × 100 × 10 − 2 × 0.1 $0.5 \times 100 \times 10^{-2} \times 0.1$, which equals 0.05. A signal-to-noise ratio of 0.05 means the signal is buried deep in the noise. It would be essentially impossible to see.

The improvement factor is the ratio of the two, 25 0.05 $\frac{25}{0.05}$, which is 500.

This numerical example powerfully demonstrates the orders-of-magnitude sensitivity gain that polarization spectroscopy provides under realistic laboratory conditions. It can turn an undetectable signal into a very clear one.

Page 64:

Let's discuss one final, but very important, practical consideration: noise mitigation. Specifically, a type of noise that arises from interference. This is Slide 25.

A common problem in these experiments is that a small amount of the strong pump light can be back-scattered from the cell windows or other optical surfaces. This scattered pump light can travel back along the probe beam's path and interfere with the probe beam at the detector.

Because the path length of this scattered light is sensitive to tiny vibrations and fluctuations in air density along the beam path, the phase of the interference term, $\phi(t)$, will fluctuate randomly. This creates a low-frequency, drifting noise component that can be very difficult to deal with.

The remedy is an elegant technique called **rapid phase modulation** or **phase dithering**. The idea is to intentionally modulate the phase of the pump beam at a high frequency. This is typically done by mounting the pump beam's retro-reflecting mirror, M2, on a piezoelectric transducer, or PZT.

By applying a sinusoidal voltage to the PZT, we can make the mirror oscillate back and forth. If we make the amplitude of this oscillation large enough, such that the optical path length changes by more than a wavelength of the light, we are rapidly sweeping the phase of the pump beam through many cycles of 2π.

Page 65:

How does this rapid phase modulation help?

The key is in how we detect our signal. We use a lock-in amplifier, and we set its integration time constant to be very long, which corresponds to a detection frequency that is much, much lower than the rapid modulation rate, f, of the piezo.

The coherent interference noise term now oscillates at this high frequency f. By averaging or integrating the signal over a time period much longer than $1/f$, this rapidly oscillating interference term averages to zero.

However, the real polarization spectroscopy signal, which arises from the pump-induced birefringence, is an *incoherent* process with respect to this phase modulation. The population changes induced by the pump depend on its intensity, not its phase. Therefore, the real signal is not affected by the phase dither.

The result is that the averaging process suppresses the unwanted coherent interference terms, while retaining our desired signal.

Page 66:

This diagram provides a schematic of the phase dithering technique.

Panel A, on the top left, shows the setup. The pump beam hits mirror M2, which is mounted on a Piezo Transducer (PZT). A function generator drives the PZT with a fast sinusoidal signal, f m o d f_{mod}, causing the mirror position, and thus the pump path length Δ L (t) $\Delta L(t)$, to oscillate rapidly.

Panel B shows the sinusoidal drive voltage applied to the PZT as a function of time.

Panel C illustrates the timing and averaging principle. The rapid phase modulation signal, f m o d f_{mod}, represents the fast-oscillating interference noise. The lock-in amplifier effectively integrates the total signal over a long time window, shown as the orange box. The key condition is that the lock-in's detection frequency, f f, is much slower than the modulation frequency, f m o d f_{mod}.

As the mechanism box explains, the fast-oscillating cosine term from the interference noise averages to zero over the long integration window, effectively eliminating this source of noise from our measurement. This is a common and essential trick for achieving the highest sensitivity in many optical experiments.

Page 67: Let's now conclude with a summary of the numerous advantages of Polarization Spectroscopy

First, **High resolution**. Like other Doppler-free techniques, the resolution is not limited by thermal motion. In practice, the primary limitation on resolution is often the small residual angle between the pump and probe beams. Any non-zero angle re-introduces a small amount of Doppler broadening, known as the residual Doppler width. With careful alignment, this can be made extremely small, allowing for exceptionally high-resolution measurements.

Second, **Superior sensitivity**. This is the key advantage we have emphasized throughout. Typically, polarization spectroscopy offers a sensitivity that is 100 to 1000 times greater than that of standard saturation spectroscopy. The ultimate limit on this sensitivity is determined by the quality of the polarizers (the extinction ratio, ξ ξ) and the quality of the cell windows (the residual window birefringence).

Page 68:

Continuing with the advantages:

Branch selectivity. As we demonstrated with the Cesium dimer example, we can choose the pump polarization—linear or circular—to selectively emphasize either P and R lines, or Q lines. This makes polarization spectroscopy a powerful diagnostic tool for making unambiguous assignments in complex, congested spectra.

Intrinsic dispersion output. By simply uncrossing the polarizers by a tiny amount, the technique naturally produces a dispersion-shaped signal. This provides an ideal error signal for laser frequency stabilization, without the need for any extra frequency dithering of the laser itself. This simplifies the experimental setup and avoids adding unwanted modulation sidebands to the laser.

Broad applicability. The principles we've discussed are very general. Polarization spectroscopy has been successfully demonstrated and widely used on a vast range of systems, including many different atoms and molecules, such as Iodine (I_2), Cesium dimer (Cs_2), and even more exotic species like rare-gas excimers.

Page 69:

Finally, to put everything together, here is a concluding roadmap for implementing polarization spectroscopy successfully in the laboratory. This is a practical checklist for any student setting out to build such an experiment.

1. First, procure the best polarizers you can afford. Ensure you are using **high extinction-ratio polarizers**, with ξ ξ value of less than $10 - 7$ 10^{-7}, if possible. This is the foundation of the technique's sensitivity.

2. Next, **characterize and, if necessary, compensate for the window birefringence**. This might involve testing the empty cell for any polarization rotation or even building a mechanical clamp to apply a compensating stress to the windows.

3. **Align your beams carefully**. Set the angle between the probe and pump beams to be as small as your optics will permit. This is crucial to minimize the residual Doppler broadening and achieve the highest possible spectral resolution.

4. Finally, to eliminate noise from back-scattered light, **use piezo modulation** on the pump beam's retro-reflector to implement the phase dithering technique we discussed.

Page 70:

And the final step on our roadmap:

Choose your pump polarization—circular versus linear—strategically, according to the desired transition branch you wish to study. If you are trying to identify Q-branch transitions in a molecule, use a linear pump. If you are interested in the P and R branches, or if you are working with an atomic transition that behaves like a P or R branch, use a circular pump. This deliberate choice is one of the most powerful features at your disposal.

By following these steps, you can harness the full power of this elegant and sensitive spectroscopic method.

That concludes our lecture on Polarization Spectroscopy. Thank you.