Vol. 2

Chapter 1.7

Page 1

Good morning, everyone. Welcome back to Physics 608, Laser Spectroscopy. I'm Distinguished Professor Dr M A Gondal, and today, we begin a new and very powerful topic, which corresponds to section 1.7 in your textbook.

Page 2:

The focus of our lecture today is on two closely related, high-resolution techniques: Laser Magnetic Resonance and Laser Stark Spectroscopy. Now, these methods represent a clever and fundamentally different approach to performing spectroscopy, a shift in strategy that becomes incredibly powerful under certain circumstances, as we are about to see.

Page 3:

Alright, let's begin with a crucial introductory question that gets to the very heart of today's topic: Why would we want to "Move the Molecule," not the Laser?

To understand this question, let's first quickly recall the paradigm of traditional laser spectroscopy, which we've been discussing in the previous sections of this course. In the standard approach, we take a tunable laser, and we systematically scan its optical frequency, which we'll call $\nu_L = \omega_L / 2\pi$. The goal, as you know, is to find a coincidence. We are searching for the specific laser frequency that exactly matches one of the fixed, inherent transition frequencies of our molecule, which we denote as

v i k ν_{ik}, representing the transition from some state 'i' to state 'k'. The molecule's energy levels are fixed, and we sweep our laser to find them.

Page 4:

This traditional approach is powerful, but it runs into a significant practical limitation. The fact is, bright, powerful, and broadly tunable laser sources simply do not exist in every single spectral window of interest. There are gaps in our technological capabilities.

A prime example, and one that is critically important for chemistry and physics, is in the mid-infrared region. Specifically, the windows from about 3 to 5 micrometers and the region around 10 micrometers are notoriously difficult. While we have some tunable sources like Quantum Cascade Lasers, or QCLs, they don't cover everything, and sometimes we need more power or stability than they can offer.

So, what do we do when we can't tune the laser to match the molecule? This brings us to the alternative strategy introduced in section one point seven. The logic is, if you can't move the mountain to Mohammed, you move Mohammed to the mountain.

The strategy is this:

- First, we take a strong, stable, *fixed-frequency* laser. We don't tune it at all. - Second, instead of tuning the laser, we tune the molecular transition itself. We literally shift the molecule's energy levels using an external field, either a magnetic field, which we denote with a capital B, or an electric field, capital E, until the transition frequency comes into resonance with our fixed laser line.

This is the core concept for the entire lecture. We are flipping the script: fixed laser, tunable molecule.

Page 5:

This alternative strategy gives rise to two powerful and complementary techniques, which form the subject of our lecture.

First, we have Laser Magnetic Resonance, which is almost always abbreviated as LMR. As the name implies, this technique uses a magnetic field to tune the molecule, and it does so by exploiting the Zeeman effect—the splitting of energy levels in a magnetic field.

Second, we have Laser Stark Spectroscopy. This is the electrical analogue to LMR. It exploits the Stark effect, which is the shifting and splitting of energy levels in an electric field.

These are two sides of the same coin, both based on the principle of tuning the molecule rather than the laser, but using different physical interactions to achieve it.

Page 6:

This slide presents a wonderful conceptual decision tree that perfectly summarizes the two main strategies in laser spectroscopy. Let's walk through it together.

At the very top, we have our ultimate goal: Achieve Resonance. That means the laser frequency, v L ν_L, must equal the molecular transition frequency, v i k ν_{ik}.

Now, we have two paths to get there.

On the left, we have "Strategy 1: Tune the Laser." The simple diagram here shows two fixed molecular energy levels, E i E_i and E k E_k, with a fixed energy gap, corresponding to the transition frequency v i k ν_{ik}. The laser, represented by the wavy blue line, is scanned in frequency to find this resonance. The goal is to find the laser frequency v L ν_L that matches the fixed molecular frequency. To do this, we use bright, broadly tunable sources like dye lasers, Titanium-Sapphire lasers, or QCLs. But, as we noted, the practical limitation is that such sources may not be available or practical in all spectral windows, especially the mid-IR. This is the content of sections 1.1 to 1.6.

Now, look at the right side: "Strategy 2: Tune the Molecule." This is our new approach. Here, the diagram shows a fixed laser frequency, v L ν_L, represented by the wavy red line. The molecular energy levels, however, are now shown with dashed lines, indicating that they are being tuned. The goal here is to tune the molecular frequency v i k ν_{ik} to match the *fixed* laser frequency v L ν_L. How do we do this? We use strong, fixed-frequency lasers, like gas lasers—Carbon Monoxide or Carbon Dioxide lasers are classic examples. The mechanism is to apply an external electric field (E E) or magnetic field (B B) to shift the molecular energy levels via the Stark or Zeeman effect, respectively. This is the subject of section 1.7, and it's our focus for today.

This side-by-side comparison is the key takeaway. It's a choice between tuning your source or tuning your sample.

Page 7

So, which spectral windows benefit the most from this new strategy of tuning the molecule? There are two key regions where this approach has proven exceptionally valuable.

First is the Mid-Infrared, or "Fingerprint" Region. This typically covers the ranges from 3 to 5 micrometers and also the region around 10 micrometers. This area of the electromagnetic spectrum is incredibly important because it's where the fundamental vibrational bands of most molecules reside. These vibrations are highly specific to a molecule's structure and bonding, acting like a unique molecular signature or "fingerprint." Observing these transitions is key to identifying and quantifying chemical species. And it just so happens that in this very region, we have a variety of intense, stable, fixed-frequency gas lasers available, such as Hydrogen Fluoride (HF), Deuterium Fluoride (DF), Carbon Monoxide (CO), Nitrous Oxide (N two O), and Carbon Dioxide (C O two) lasers. The existence of these powerful lasers makes this region a perfect candidate for LMR and Stark spectroscopy.

The second key area is the Far-Infrared, which is also called the Terahertz region. This corresponds to wavelengths from roughly 30 micrometers up to 1000 micrometers, or one millimeter. This is a lower energy region where the pure rotational lines of polar molecules reside. Studying these transitions gives us exquisitely precise information about molecular geometries, bond lengths, and angles. And just like in the mid-IR, this region has excellent fixed-frequency laser sources that we can leverage.

Page 8:

Continuing with the sources available in the far-infrared, we have some very important examples of bright, fixed-frequency lasers.

There are water vapor, H 2 O H_2O, and heavy water, D 2 O D_2O, gas lasers, which have strong emission lines around a wavelength of 125 μ m $125\,\mu\text{m}$. There are also Hydrogen Cyanide, HCN, lasers, which operate around 330 μ m $330\,\mu\text{m}$.

Furthermore, the list of available fixed-frequency lines has been greatly expanded by the development of optically pumped molecular lasers. In these devices, a powerful laser, like a C O 2 CO_2 laser, is used to "pump" a different gas, which then lases on a far-infrared transition. This technique has generated a list of over one thousand distinct far-infrared laser lines, providing a dense, albeit not continuous, coverage of the spectrum.

So, why are these fixed-frequency lasers so attractive? What makes them worth designing an entire spectroscopic technique around? There are three main reasons.

First, they offer very high output power. We're often talking about Watts of power, not the milliwatts that are typical for many tunable sources. High power is crucial for sensitivity.

Second, they have excellent spectral purity and long-term frequency stability. These gas lasers can be engineered to produce an extremely narrow, stable output frequency, which is essential for high-resolution measurements.

And third...

Page 9:

...is their commercial availability and relative simplicity compared to many broadly tunable sources. Building and operating a stable, powerful gas laser can often be a more robust and cost-effective solution than implementing a complex tunable system, especially for these challenging mid- and far-infrared spectral regions. These are workhorse lasers that have been staples of research labs for decades for good reason.

Page 10:

This graph provides a powerful visual summary of the points we've just discussed. It plots the "Power versus Wavelength for Common Fixed-Frequency Lasers."

Let's break it down. The vertical axis is the typical output power in Watts. The horizontal axis is the wavelength, λ, in micrometers, plotted on a logarithmic scale, spanning from 1 micrometer to 1000 micrometers.

The plot is divided into two regions. On the left, shaded in blue, is the "Mid-IR 'Fingerprint'" region. Here you can see several vertical bars representing specific gas lasers. We have the HF laser around 3 micrometers, producing about 3 watts. The DF laser is next, slightly longer in wavelength, with about 2.5 watts. Then the CO laser, near 5 micrometers, putting out 4 watts. And look at the enormous bar for the CO_2 laser, around 10 micrometers, reaching up to 8 watts! This is a tremendous amount of power. There's also a smaller bar for N_2 O.

Now, look at the region shaded in orange on the right: the "Far-Infrared / Terahertz" region. Here we see the bars for the H_2O laser around 125 micrometers, the D_2O laser, and the HCN laser near 300 micrometers. Notice their power is lower, in the range of a few hundred milliwatts to maybe a watt, but they are still very important sources in this difficult spectral region.

Most importantly, I want you to notice the annotations labeled "Gap in Tunability." There are large empty spaces between these powerful fixed lines. This graph visually demonstrates the problem: if your molecule's absorption doesn't happen to coincide with one of these laser lines, traditional spectroscopy is impossible. This is the gap that LMR and Stark spectroscopy are designed to bridge.

Finally, notice the dashed line on the far right, labeled "Optically Pumped Molecular Lasers." It shows that these sources provide a dense set of lines, but with lower power, helping to "partially fill" the far-infrared gap, giving us more fixed-frequency options to work with.

Page 11:

Alright, now that we understand the motivation, let's dive into the physics of the first technique: Laser Magnetic Resonance. We'll start with the fundamentals of molecular energy levels in a magnetic field.

Any given molecular energy level is labeled by a set of quantum numbers. For our purposes, the most important ones are: First, the vibrational quantum number, which we denote with a lowercase ‘v’. Second, the rotational quantum number, capital ‘J’. It is crucial to remember that J J

represents the total *mechanical* angular momentum of the molecule. This means it includes contributions from electron spin and orbital angular momentum, as well as molecular rotation, but it *excludes* the nuclear spin angular momentum.

Now, for today's topic, the star of the show is the magnetic quantum number, which we denote with a capital 'M', or sometimes M J M_J. This quantum number represents the projection of the total angular momentum vector, J J, onto a defined axis, which we will choose to be the axis of the external magnetic field. The possible values of M M range in integer steps from − J $-J$ to + J $+J$.

The final point on this slide is the critical starting point for our entire discussion: In zero magnetic field, all of these 'M' sublevels are degenerate. That is, for a given J J, all 2 J + 1 $2J+1$ possible orientations of the angular momentum have the exact same energy, which we'll call E 0 E_0.

Page 12:

So, what happens when we disturb this situation? We apply a homogeneous, static magnetic field. Let's define our coordinate system such that the field points along the z-axis, so the vector B $\mathbf{B}$ is equal to the magnitude B B times z ^ $\hat{\mathbf{z}}$, the unit vector in the z-direction.

The immediate consequence is that the degeneracy we just mentioned is lifted. The single energy level splits into 2 J + 1 $2J+1$ distinct components, which are known as the Zeeman components.

What is the physical origin of this splitting? It arises from the interaction of the molecule's own magnetic dipole moment, μ μ, with the external magnetic field, B $\mathbf{B}$. This is a potential energy interaction, just like a compass needle aligning in the Earth's magnetic field.

For an atom or a molecule, the magnetic dipole moment is fundamentally linked to its angular momentum. The magnetic dipole moment vector, μ μ, is proportional to the total angular momentum vector, J J. The quantum mechanical relationship is given by the equation you see here:

μ = g μ 0 J $\hbar$.

$$\boldsymbol{\mu} = g\,\mu_0\,\frac{\mathbf{J}}{\hbar}.$$

- g

is the Landé g-factor, a dimensionless proportionality constant that we will discuss in a moment. - μ ⊂ z e r o

(or μ $_0$

) is a fundamental constant called the Bohr magneton. It is defined as 'e' times $\hbar$ $\hbar$ divided by two times the electron mass m e m_e, where 'e' is the elementary charge. Its value is approximately nine point two seven four times ten to the minus twenty-four Joules per Tesla. The Bohr magneton is the natural unit for magnetic moments that arise from electron orbital or spin angular momentum.

So this equation tells us that the magnetic moment is simply a scaled version of the angular momentum.

Page 13:

Let's take a closer look at the Landé g-factor, denoted by 'g'.

This is a dimensionless factor that essentially tells us how strong the magnetic moment is for a given amount of angular momentum. Its precise value is very important as it depends intimately on the molecule's electronic configuration and the specific way in which the various internal angular momenta (like electron spin and orbital angular momentum) couple together. This is described by formalisms like Hund's coupling cases, which you may have encountered in quantum mechanics.

For many of the species we study with LMR, particularly radicals which have unpaired electrons, the g-factor is often close to 2, which is the value for a free electron spin. However, the deviation from this value is unique to the molecule and its quantum state, making the experimental determination of the g-factor a powerful tool for characterizing the molecule.

Page 14: Zeeman Splitting of a Rotational Level

This diagram provides a clear and simple illustration of Zeeman splitting. The title is "Zeeman Splitting of a Rotational Level" for the specific case where the total angular momentum quantum number, J J, is equal to 2.

Let's look at the axes. The vertical axis represents energy. The horizontal axis represents the strength of the applied magnetic field, B B.

On the left side, where the magnetic field B B is equal to zero, we see a single, thick horizontal line. This represents the single, degenerate energy level for J = 2 $J = 2$. All possible orientations of the angular momentum have the same energy.

Now, follow the faint lines to the right side of the diagram, where the magnetic field B B is greater than zero. That single level has now split into 2 J + 1 $2J + 1$, which is 2 × 2 + 1 $2 \times 2 + 1$, or 5, distinct and equally spaced energy levels. Each of these sublevels corresponds to a different magnetic quantum number, M M. You can see them labeled from top to bottom: M = + 2 $M = +2$, M = + 1 $M = +1$, M = 0 $M = 0$, M = − 1 $M = -1$, and M = − 2 $M = -2$.

This picture is the essence of the Zeeman effect: an external magnetic field breaks the spatial degeneracy and provides us with a set of levels whose energy we can now tune.

Page 15:

Now let's move from the qualitative picture to a quantitative, step-by-step derivation of the Zeeman energy shift.

We begin with the classical expression for the potential energy, U U, of a magnetic dipole μ μ in a magnetic field B B. This is given by the dot product:

U = − μ ⋅ B .

$$U = -\mu \cdot B.$$

To simplify this, as we did before, we choose our coordinate system so that the magnetic field B B is aligned purely along the z z-axis. This reduces the dot product to a simple scalar multiplication: U U equals negative μ z μ_z times B B, where μ z μ_z is the z z-component of the magnetic dipole moment.

Now, we need the quantum mechanical value for this energy. The energy shift for a particular sublevel is the expectation value of this interaction Hamiltonian. The slide shows the result of substituting our operator for μ z μ_z. The expectation value of the z z-component of the magnetic moment, for a state described by quantum numbers J J and M M, is given by g g times the Bohr magneton μ 0 μ_0 times M M.

This result comes directly from fundamental quantum mechanics. The operator for the z z-component of the magnetic moment, μ z μ_z, is

μ z = g μ 0 J z ħ .

$$\mu_z = \frac{g\,\mu_0 J_z}{\hbar}.$$

The states | J , M ⟩ $|J, M\rangle$ are eigenstates of the operator J z J_z, with the well-known eigenvalue equation

J z | J , M ⟩ = ħ M | J , M ⟩ .

$$J_z|J, M\rangle = \hbar\, M|J, M\rangle.$$

When you calculate the expectation value, the ħ ħ terms cancel, and you are left with the simple result that the energy shift is directly proportional to the magnetic quantum number M M.

Page 16:

This leads us directly to the final, crucial expression for the energy of a Zeeman-shifted level. The energy, E E, which is a function of the quantum number M M and the field strength B B, is given by:

E (M , B) = E 0 − g μ 0 B M .

$$E(M,B) = E_0 - g\mu_0 BM.$$

Let's recap the key variables here to make sure we're all on the same page: - E 0 E_0 is the original energy of the level in zero magnetic field. - B B is the magnitude of the applied magnetic field, measured in tesla. - The term g μ 0 B M $g\mu_0 BM$ represents the energy shift itself.

It is very important to pay attention to the sign convention. Notice the negative sign in the equation. For a molecule with a positive g-factor, which is common, this means that a sublevel with a positive M M quantum number will be shifted down in energy. Conversely, a negative M M value will be shifted up in energy. The M = 0 $M = 0$ sublevel, of course, experiences no shift from this term.

So, how large is this shift in practice? Let's do a quick, typical magnitude estimate.

Page 17:

Let's plug in some realistic numbers. We'll consider a molecule where the g-factor is approximately 2 2, a typical value for a radical. Let's assume we apply a strong magnetic field from a laboratory electromagnet, say B = 2

$B = 2$ Tesla. And let's look at a sublevel where the magnitude of M M is equal to 1 1.

The magnitude of the energy shift, absolute value of Δ E ΔE, is then approximately:

| Δ E | ≈ 2 × (9.274 × 10 − 24 J / T) × 2 T .

$$|\Delta E| \approx 2 \times (9.274 \times 10^{-24}\,\mathrm{J/T}) \times 2\,\mathrm{T}.$$

This gives an energy shift of approximately 3.7 × 10 − 23 3.7×10^{-23} Joules.

Now, an energy in Joules isn't always the most intuitive unit for a spectroscopist. A more common unit is the wavenumber, or inverse centimeters. If we convert this energy shift into wavenumbers, by dividing by h h and c c, it corresponds to a wavenumber shift, Δ v ~ $\Delta\tilde{\nu}$, of approximately 2 c m − 1 $2\,\mathrm{cm}^{-1}$.

This is a critically important result. Two inverse centimeters is a very significant and easily resolvable shift.

For comparison, the Doppler width of a line in the infrared might be around 0.01 c m − 1 $0.01\,\mathrm{cm}^{-1}$. So a 2 c m − 1 $2\,\mathrm{cm}^{-1}$ tuning range is enormous! This calculation proves that using laboratory-scale magnetic fields, we can tune molecular transitions by a substantial amount, making this technique highly practical.

Page 18: Energy Level Splitting in a Magnetic Field (Zeeman Effect)

Here we have another excellent diagram visualizing the Zeeman effect, this time plotting the energy of the sublevels as a function of the magnetic field strength.

The title is "Energy Level Splitting in a Magnetic Field (Zeeman Effect)."

The vertical axis is Energy, E E, and the horizontal axis is the Magnetic Field, B B.

At the far left, where the magnetic field is zero, all the sublevels start at the same energy, E 0 E_0. As we move to the right and increase the magnetic field, the levels split apart. This graph shows the splitting for a J = 1 $J = 1$ level, which has M M sublevels of + 1 $+1$, 0 0, and − 1 -1.

The line for M = 0 $M = 0$ is a flat, horizontal gray line. Its energy does not change with the magnetic field.

The line for M = + 1 $M = +1$, shown in red, slopes downwards. Its energy decreases linearly as B B increases.

The line for M = − 1 $M = -1$, shown in blue, slopes upwards. Its energy increases linearly with B B.

This linear dependence is exactly what our equation predicted: The energy shift, Δ E ΔE, which is E − E 0 $E - E_0$, is equal to − g μ B $-g\mu_B$ times B B times M M.

Δ E = − g μ B B M

$$\Delta E = -g\,\mu_B\,B\,M$$

This equation is written at the top of the graph.

The diagram also correctly labels the energy separation between the M = 0 $M = 0$ and M = + 1 $M = +1$ levels at a given field B B as g μ B B $g\,\mu_B\,B$.

This plot provides a perfect dynamic picture of how we "tune the molecule" with a magnetic field.

Page 19

Now that we understand how individual energy levels shift, we can consider what happens to the frequency of a *transition* between two such levels in a magnetic field.

We are interested in a spectroscopic transition from a lower state, which we'll label with quantum numbers (v ′ ′ , J ′ ′ , M ′ ′) (v'', J'', M''), to an upper state, labeled (v ′ , J ′ , M ′) (v', J', M'). The single prime denotes the upper state and the double prime denotes the lower state, by standard spectroscopic convention.

In the absence of a magnetic field, the optical angular frequency of this transition, ω 0 ω_0, is given by the Bohr frequency condition:

ω 0 = E 0 ′ − E 0 ′ ′ ℏ .

$$\omega_0 = \frac{E_0' - E_0''}{\hbar}.$$

To find the new, field-dependent frequency, we simply calculate the energy of the upper level in the field and the energy of the lower level in the field, and subtract them. The new transition frequency is obtained by subtracting the Zeeman shifts of the upper and lower levels from their respective zero-field energies before taking the difference.

Page 20:

This leads us to the final equation for the transition frequency, ω ω, as a function of the magnetic field, B B:

ω (B) = ω 0 − μ 0 ℏ (g ′ M ′ − g ″ M ″) B .

$$\omega(B) = \omega_0 - \frac{\mu_0}{\hbar}(g'M' - g''M'')\,B.$$

Let's dissect this. The new frequency, ω (B) $\omega(B)$, is the zero-field frequency, ω 0 ω_0, plus a shift term. This shift term depends linearly on the magnetic field B B. Crucially, it also depends on the g-factors and M-quantum numbers of both the upper state (g ′ g', M ′ M') and the lower state (g ″ g'', M ″ M'').

Now, not all possible transitions between M-sublevels are allowed. For electric-dipole transitions in a magnetic field, we have selection rules. The selection rule for the magnetic quantum number M M is that Δ M ΔM, which the slide defines as M ″ − M ′ $M'' - M'$, must be equal to 0 0 or ± 1 ± 1. The selection rule for J J is the standard rotational rule: Δ J ΔJ can be 0 0 or ± 1 ± 1, with the exception that a transition from J = 0 $J = 0$ to J = 0 $J = 0$ is forbidden.

The resulting spectral pattern, therefore, consists of up to three distinct groups of transitions, often labeled with Greek letters:

- π π components correspond to Δ M = 0 $\Delta M = 0$. - σ + σ^+ and σ − σ^- components correspond to Δ M = + 1 $\Delta M = +1$ and − 1 -1 respectively.

There's a special case known as the "normal" Zeeman effect, which occurs if the g-factors of the upper and lower states are identical, so g′ = g″ $g' = g''$. In this case, the complex pattern collapses to a simple, symmetric triplet of lines. However, in molecules, g′ g' and g″ g'' are almost never equal, leading to what is called the "anomalous" Zeeman effect, which is actually the norm for molecules.

Page 21

So, what determines the practical tuning range of our molecular transition? Looking back at the frequency shift equation, we can see that the tuning range scales with two key parameters. First, the strength of the magnetic field, B B, that we can apply. And second, the magnitude of the difference between the g-factors of the two states, that is, the absolute value of |g″ − g′| $|g'' - g'|$.

This second point has a very important practical implication: Molecules that have a large difference in their g-factors between the upper and lower states are the best candidates for LMR spectroscopy. This is because their transition frequencies will tune much more rapidly with the applied magnetic field, making it easier to bring them into resonance with a fixed-frequency laser.

Molecules with unpaired electrons, known as radicals, are perfect examples. The unpaired electron spin contributes significantly to the g-factor, and this contribution can be different in the two vibrational or rotational states involved in the transition, leading to a large g-factor difference and making them ideal targets for LMR.

And just to remind ourselves of the scale, a magnetic field B B of up to 2 T $2\,\mathrm{T}$ can produce shifts on the order of 2 c m − 1 $2\,\mathrm{cm}^{-1}$, providing a substantial tuning range.

Page 22:

Now let's put all these pieces together into a conceptual picture of an LMR scan. The experimental procedure is elegant in its simplicity.

First, you fix the laser at a single, stable frequency, ω L ω_L. You do not change it.

Second, you place your sample in a magnetic field, and you sweep the strength of this field, B B, linearly in time, for example from 0 up to 2 Tesla.

Third, as you sweep the field, the frequencies of the various Zeeman components of your molecular transition are tuning, each at its own rate. At some point, the frequency of one of these Zeeman components will cross the fixed laser frequency, ω L ω_L. This resonance condition occurs at a unique magnetic field value, which we'll call B r e s B_{res}, satisfying the equation:

ω (B r e s) = ω L .

$$\omega(B_{res}) = \omega_L.$$

Fourth, when this resonance occurs, the molecules in the sample absorb the laser light. This absorption is detected as a signal, typically a peak or a dip in the measured laser output intensity. The spectrum you record is therefore a plot of this signal versus the magnetic field, B B.

Page 23:

This leads to a very important point of clarification. The final spectrum we obtain from an LMR experiment is conceptually comparable to a traditional "frequency domain" spectrum where you plot signal versus frequency. However, in LMR, the horizontal axis is now the magnetic field strength, B B, instead of the frequency, ν ν.

We are essentially using the magnetic field as a proxy variable to scan through the molecular energy landscape. The position of a peak in the magnetic field domain gives us precise information about the molecular transition frequency, but the experiment itself is performed by sweeping a field, not a frequency.

Page 24:

This slide provides a textual description of the conceptual diagram we're about to see, which illustrates the fundamental principle of Laser Magnetic Resonance spectroscopy. I'll read through it as it sets the stage perfectly.

The diagram shows how the frequencies of molecular transitions, ω ω, change as a magnetic field, B B, is applied. Due to the Zeeman effect, a single transition that exists at zero-field, which we call ω 0 ω_0, splits into multiple components. The diagram shows three representative groups: σ − σ^-, for Δ M = − 1 $\Delta M = -1$; π π, for Δ M = 0 $\Delta M = 0$; and σ + σ^+, for Δ M = + 1 $\Delta M = +1$. Each of these groups tunes at a different rate as the field is applied, which is represented by the slope of the lines on the graph.

In an LMR experiment, the laser frequency, ω L ω_L, is held constant while the magnetic field is swept. A resonance, and therefore an absorption signal, is detected each time a molecular transition's frequency tunes into coincidence with the laser frequency. The bottom part of the diagram shows the resulting absorption spectrum, where peaks appear at the specific magnetic field values where these resonances occur.

The key takeaway is that this technique effectively converts a frequency-domain problem—finding a resonance frequency—into a magnetic field-domain measurement.

Page 25:

Now, let's look at the diagram itself. This is a beautiful illustration of the entire LMR process.

Let's start with the top panel, which plots Frequency, ω ω, on the y-axis versus Magnetic Field, B B, on the x-axis. At B = 0 $B = 0$, on the far left, you see a single point labeled ω 0 ω_0. This is our single, unsplit molecular transition. As we increase the magnetic field, this single transition splits into three lines, each with a different slope, representing the different tuning rates. These are our σ − σ_-, π π, and σ + σ_+ components. The blue line has the steepest positive slope, the green line has a shallower slope, and the red line actually has a negative slope in this example.

Now, look at the dashed pink line that runs horizontally across the graph. This represents our fixed laser frequency, ω L ω_L. The experiment consists of finding where the tuning molecular lines intersect this fixed laser line. You can see three distinct intersection points, marked with solid circles.

The blue line intersects at a low magnetic field. The green line intersects at a medium field. And the red line intersects at a high magnetic field. These are our three resonance conditions.

Now, let's look at the bottom panel. This plots the Absorption Signal versus the Magnetic Field, B B. Vertically aligned below each intersection point in the top panel, you see a corresponding absorption peak in the bottom panel. The first resonance gives the blue peak labeled σ − σ_-. The second gives the green peak labeled π π. And the third gives the red peak labeled σ + σ_+.

This two-panel diagram perfectly shows how a scan in magnetic field produces a spectrum that resolves the individual Zeeman components of the molecular transition.

Page 26:

Having understood the concept, let's now look at a typical experimental setup for intracavity LMR. The term "intracavity" is key here.

First, we have a laser resonator, which is the heart of the laser, formed by two mirrors, let's call them A and B. For far-infrared LMR, this might be an H two O water vapor laser, for example.

The second and most crucial feature of this setup is that the sample cell, containing the gas we want to study, is placed *inside* the laser cavity, between the two mirrors. Why would we do this? The reason is to exploit the incredibly high photon flux that exists inside a laser resonator. The circulating power inside the cavity can be easily 100 times, or even more,

than the power that actually comes out of the laser. By putting the sample where the light is most intense, we dramatically increase the interaction strength and therefore the sensitivity of our measurement. This allows us to detect very small concentrations of molecules.

Third, for studying unstable species like radicals, the sample cell is often incorporated into a continuous-flow system. A constant flow of gas passes through the cell, allowing us to maintain a steady-state concentration of the short-lived molecules we want to study.

Page 27:

Let's continue with some more details of the intracavity setup.

How do we generate the radicals we want to study? They are typically created right inside the apparatus, very close to the laser beam. Common methods include using a microwave discharge to break apart stable precursor molecules, or using chemical addition, where a reactant gas is mixed in to produce the desired radical through a chemical reaction.

Now, a very clever component in many far-infrared LMR spectrometers is a polyethylene membrane that serves as a beam splitter. This thin plastic film has two critical functions. First, it physically separates the laser gain medium, like the water vapor, from the sample region where our radicals are. This is important to prevent the sample from interfering with the laser's operation.

Second, and perhaps more ingeniously, the membrane is often placed at the Brewster angle, which means it also acts as a partial polarizer. This allows us to selectively observe different types of Zeeman transitions. By

rotating the entire sample tube assembly around the optical axis, we can align the electric field of the laser radiation either parallel or perpendicular to the magnetic field. This enables us to select either the $\Delta M = 0$ (π) transitions or the $\Delta M = \pm 1$ (σ) transitions. This ability to select transitions is a huge help when it comes to assigning and interpreting a complex spectrum.

Finally, the detection scheme itself is quite straightforward, as we'll see on the next slide.

Page 28:

The detection scheme for LMR is conceptually simple. You place a photodetector outside the laser's output mirror to monitor the laser's output intensity. Then, you simply record the signal from this detector as you scan the magnetic field, B.

When the magnetic field tunes a molecular transition into resonance with the fixed laser frequency, the molecules inside the cavity absorb photons. This absorption introduces a loss into the laser cavity. Because the laser's output power is very sensitive to intracavity losses, this absorption causes a small but detectable change—usually a dip—in the output intensity.

The result is that you obtain resonance curves, which look like absorption dips or sometimes emission spikes depending on the specific laser dynamics and experimental configuration, plotted as a function of the scanned magnetic field.

Page 29:

This diagram shows a complete, typical intracavity LMR experimental setup, putting all the pieces we've just discussed together. Let's trace the path.

On the far left and far right, we have the Resonator Mirrors, A and B, which form the optical cavity. The laser beam travels back and forth between them.

The cavity is divided into two sections. On the left is the Laser Gain Medium, for example, water vapor, which is introduced through the gas inlet. This is where the laser light is generated.

On the right is the Sample Cell, which is part of a continuous flow system, with gas coming in and being pumped out. This is where our target molecules are. You can see a section labeled "Radical Generation," fed by a microwave generator, where the unstable species are created.

The two sections are separated by the Polyethylene Membrane, which is labeled as both a Beam Splitter and a Polarizer.

The entire sample cell is surrounded by a large solenoid, which generates the magnetic field, B. This solenoid is powered by a variable B-field power supply, which is responsible for scanning the field.

A note points out that the whole assembly can be rotated about the optical axis to select for pi or sigma transitions.

Finally, a small fraction of the laser light passes through the output coupler, Mirror B, and goes to a photodetector. The signal from the detector is then sent to a recorder or computer.

In the bottom right corner, you see a graph of a "Typical Resonance Signal": the laser intensity dips sharply at a specific magnetic field value, indicating a resonance. This entire setup is a beautiful example of clever experimental design to achieve very high sensitivity.

Page 30:

The intracavity technique is already very sensitive, but we can push it even further. This slide discusses techniques for boosting sensitivity, namely Field Modulation and Zero-Field LMR.

The most common technique is magnetic field modulation. In addition to the large, slowly swept DC magnetic field, we superimpose a small AC ripple, which we'll call Δ B ΔB. We modulate the field at a specific audio frequency, say a few kilohertz.

What does this do? When combined with a lock-in amplifier, this modulation technique has the effect of converting the absorption line shape into something that looks like its first derivative. You no longer see a simple dip; you see a characteristic "S-shaped" curve.

The real power comes from using a lock-in amplifier, which is synchronized to the modulation frequency. A lock-in is a phase-sensitive detector. It is exquisitely sensitive to signals that are happening at the exact reference frequency and phase, and it aggressively rejects all other noise— 1 f $\frac{1}{f}$ noise, detector noise, laser fluctuations—that is outside a very narrow bandwidth around the modulation frequency. This dramatically suppresses the noise and can improve the signal-to-noise ratio by orders of magnitude.

This method allows for truly incredible achievable detection limits.

Page 31:

Just how sensitive is this? With field modulation and lock-in detection, LMR has been used to detect the CH radical at concentrations down to 2 × 10 8 2×10^8 molecules per cubic centimeter, with a 1 s 1 s time constant. That is an astonishingly low concentration, equivalent to detecting a puff of smoke in a large concert hall. It showcases the immense power of this technique for studying trace species.

Now, let's briefly touch on a related but different concept: Zero-field LMR. In this variant, you actually use a tunable laser, and you set its frequency exactly to the center of the molecular transition at zero magnetic field, v 0 ν_0.

Then, instead of sweeping the field over a large range, you modulate the magnetic field B B symmetrically *about zero*. So the field oscillates, for instance, from − 10 Gauss -10 Gauss to + 10 Gauss $+10$ Gauss.

The clever part is this: the Δ M = + 1 $\Delta M = +1$ and Δ M = − 1 $\Delta M = -1$ Zeeman components tune in opposite energy directions as the field goes from negative to positive. This means that the signals detected by the lock-in amplifier for these two types of transitions will have opposite phases—they will be 180 ∘ 180° out of phase with each other. This provides a very clear, background-free way to discriminate between the different Zeeman components.

Page 32:

This elegant technique of Zero-Field LMR is not just a theoretical concept; it has been practically demonstrated.

A key example is the work by Urban and colleagues, who used this method to study the nitric oxide radical, NO. For their experiment, they used a spin-flip Raman laser, which is a type of tunable infrared laser, to perform zero-field LMR and clearly demonstrate the phase discrimination between the different Zeeman components.

Page 33:

This diagram visually explains the concept of Zero-Field LMR and the phase relationship in the detected signal. The title is "Zero-Field Laser Magnetic Resonance (LMR): Lock-in Amplifier Output Phase vs. Magnetic Field Modulation."

The horizontal axis represents time, or more precisely, the phase of the modulation, ω t ωt. The vertical axis is amplitude in arbitrary units.

The dashed blue curve represents the magnetic field modulation itself. You can see it's a simple sine wave, oscillating symmetrically about zero.

Now, look at the two solid curves, which represent the output signal from the lock-in amplifier. The red curve and the purple curve are the signals corresponding to the Δ M = + 1 $\Delta M = +1$ and Δ M = − 1 $\Delta M = -1$ transitions. As you can clearly see, they are also sinusoidal, but they are perfectly out of phase with each other. When the red curve is at its positive maximum, the purple curve is at its negative maximum.

The annotation with the vertical arrow explicitly points this out: "Opposite Phase (180 degree shift)." This is the signature of zero-field LMR. By setting the lock-in amplifier to detect one phase, you can selectively measure one type of transition while rejecting the other, providing an excellent method for signal discrimination.

Page 34:

Let's now look at a real-world example spectrum to see what LMR can do. This case study is of the CH radical, detected in a flame.

The source of the molecules is a low-pressure flame, created by burning acetylene, C 2 H 2 C_2H_2, with oxygen, O 2 O_2. Flames are complex, high-temperature environments, making them a challenging target for spectroscopy.

The laser used was a fixed-frequency water vapor, H 2 O H_2O, far-infrared laser line.

The observed LMR spectrum, which is a plot of signal versus magnetic field, contains two main features.

First, we see a set of sharp resonances that belong to our target molecule, the CH radical.

Second, we see some broader, overlapping lines. These are due to the OH radical, which is a very common background radical that is almost always present in flames.

The interpretation of such a spectrum requires us to disentangle these contributions.

Page 35:

So, how do we interpret a spectrum like this? The information we can extract is incredibly rich.

First, the *positions* of the CH peaks—the exact magnetic field values at which they occur—are the primary data. By analyzing these positions, and knowing the laser frequency precisely, we can work backwards to determine fundamental molecular parameters. This includes precise values for the rotational constants and, most importantly for LMR, the magnetic parameters, namely the Landé g-factors for the upper and lower states.

Second, the *relative intensities* of the different peaks also contain valuable information. The intensity of each transition is proportional to the population of the initial M sublevel. In a sample at thermal equilibrium, these populations follow a Boltzmann distribution. Therefore, by measuring the relative intensities, we can determine the rotational temperature of the molecules in the flame.

Page 36: Example LMR Spectrum: CH Radical in a Flame

This slide presents the full figure for our case study: "Example LMR Spectrum: CH Radical in a Flame."

Let's first look at the main plot. The vertical axis is the LMR signal, and the horizontal axis is the magnetic field in units of kiloGauss, or kG, from 0 to 12. You can see a series of sharp, derivative-shaped signals. Several are

labeled "CH," indicating they belong to our target radical. There is also a broader, underlying feature labeled "OH," which is the background species.

Now let's look at the information provided in the boxes on the right, which is essential for understanding the experiment.

The "Experimental Conditions" box reminds us of the source (a low-pressure acetylene flame) and the laser (a fixed-frequency H2O FIR laser).

The "Peak Assignment" box gives an illustrative example of how the analysis is done. It shows a table. For a single rotational transition of CH, the Zeeman effect splits it into multiple lines. The table shows that the peak observed at a magnetic field of 4.2 kiloGauss can be assigned to the transition from the lower state M J ″ M_J'' of − 1 2 $-\frac{1}{2}$ to the upper state M J ′ M_J' of − 3 2 $-\frac{3}{2}$. Similarly, other peaks at 6.8 and 9.3 kiloGauss are assigned to different Zeeman transitions. This is the process of spectral assignment.

Finally, the "Interpretation" box summarizes our takeaways. Peak Positions give us molecular parameters like rotational constants and g-factors. Peak Intensities reflect the Boltzmann populations of the M sublevels. And the presence of Overlapping Lines from OH illustrates the need for high-resolution techniques like LMR to distinguish different species in a complex mixture like a flame.

Page 37:

So far, we have focused on intracavity LMR, where the sample is inside the laser. However, this is not always practical. For example, if the laser is very

large or if the sample conditions (like high pressure or corrosiveness) are incompatible with the laser optics, we need an alternative. This leads us to External-Cell LMR.

The motivation is to handle situations where the intracavity arrangement is impractical.

In an external-cell setup, how do we detect the weak absorption? A very clever way is to use polarization effects. The basic arrangement is to place the sample cell between two crossed polarizers. The first polarizer sets the polarization of the incoming light, and the second polarizer, called the analyzer, is set to be perpendicular to the first. In this configuration, if nothing happens to the light in the sample cell, no light can pass through the analyzer to the detector. This provides a "zero-background" measurement, which is excellent for sensitivity.

There are two main geometries for this. The first is the longitudinal configuration, also known as the Faraday effect configuration. In this setup, the magnetic field B B is applied parallel to the direction of light propagation.

Page 38:

In this longitudinal, or Faraday, configuration, the Zeeman resonance induces a phenomenon called circular birefringence in the sample. This means that at resonance, the refractive index of the gas becomes different for left-circularly polarized light and right-circularly polarized light.

Since linearly polarized light can be described as a superposition of left- and right-circularly polarized components, this difference in refractive

indices causes the plane of polarization of the linear light to rotate as it passes through the sample. This rotation is the Faraday effect. Now that the polarization has been rotated, it is no longer perfectly perpendicular to the analyzer, so a component of the light can pass through and generate a signal at the detector.

The second geometry is the transverse configuration, which relies on the Voigt effect. Here, the magnetic field B B is applied perpendicular to the direction of light propagation. In this case, the Zeeman resonance induces linear birefringence, meaning the refractive index is different for light polarized parallel to the B B field versus perpendicular to it. This also causes a change in the polarization state of the incident light—for example, if the incident polarization is at 45 ∘ 45° to the B B field, the outgoing light will be elliptically polarized. This altered polarization state will have a component that can pass through the crossed analyzer, again generating a signal.

In both of these external-cell methods, the sensitivity can be further boosted by using field modulation combined with lock-in detection, just as we did for the intracavity technique.

Page 39:

This diagram illustrates the two external-cell LMR configurations.

Let's start with panel (a), the Longitudinal or Faraday Configuration. We see a laser on the left, followed by a polarizer that creates a vertically polarized E E-field. The light then enters the sample cell, which is wrapped in a solenoid, indicating that the magnetic field B B is along the axis of light

propagation. Inside the cell, at resonance, the polarization plane rotates. The diagram shows the decomposition of the E E-field into left- and right-handed circular components, E − E_- and E + E_+, which experience different phase shifts, causing the rotation. After the cell, the light hits the analyzer, which is crossed (set to horizontal). Because the polarization has been rotated, a signal is transmitted to the detector.

Now look at panel (b), the Transverse or Voigt Configuration. Again, we have a laser and a polarizer. This time, the polarizer is set to create a field E E at 45 ∘ $45°$. The sample cell is placed between the north and south poles of a magnet, so the B B field is vertical, perpendicular to the light path. The incident E E-field is shown decomposed into components parallel and perpendicular to B B. These two components travel at different speeds, causing the light to become elliptically polarized as it exits the cell. This elliptical light has a vertical component that can pass through the analyzer (which is set to pass light at 135 ∘ $135°$, crossed with the initial 45 ∘ $45°$ polarization), generating a signal at the detector.

These are two elegant ways to perform LMR outside the laser cavity.

Page 40:

The techniques we've discussed so far offer high resolution, but their accuracy can still be limited by the Doppler broadening of the transitions in the gas sample. Can we do even better? The answer is yes, by combining LMR with the powerful technique of saturation spectroscopy to achieve Doppler-Free LMR.

Here's the basic idea of saturation spectroscopy. You take your fixed-frequency laser and split the beam into two: a strong "pump" beam and a weak "probe" beam. These two beams are directed to pass through the same sample volume, but in opposite, counter-propagating directions.

The intense pump beam is strong enough to saturate the transition. This means it interacts preferentially with molecules that are standing still with respect to the light—that is, those with a velocity component of zero along the laser axis. The pump beam excites these molecules to the upper state, depleting the population of the ground state for this specific velocity group. This is often called "burning a hole" in the Doppler-broadened velocity distribution.

Now, when the weak probe beam, traveling in the opposite direction, passes through the sample, it also interacts with this same zero-velocity group. However, since the pump beam has already removed many of the absorbing molecules from the ground state, the probe beam sees a reduced absorption. This sharp decrease in absorption at the exact line center is known as a Lamb dip.

The crucial step for Doppler-Free LMR is to perform this saturation spectroscopy experiment while simultaneously applying a magnetic field. This allows us to resolve the individual Zeeman components as sharp Lamb dips, completely free from the blurring effect of Doppler broadening.

Page 41:

What are the benefits of this sophisticated Doppler-Free LMR technique? They are truly significant.

First, the resolution is dramatically improved. The observed linewidth is no longer limited by the Doppler width, but rather by the much narrower natural or collisional width of the transition. This can be less than one megahertz, which is often a factor of 10 to 100 times sharper than the Doppler-limited spectrum.

Second, this exquisite resolution allows for the accurate measurement of very subtle effects. For example, we can precisely measure collisional broadening parameters—how the linewidth changes with pressure—and small pressure-induced shifts in the line position.

Third, the shape and depth of the saturation signal, the Lamb dip, are directly related to the transition strength. This allows for the direct determination of fundamental molecular properties like transition dipole moments from the analysis of the saturation parameters. This is a very powerful capability.

Page 42:

This four-panel diagram provides a comprehensive overview of the Doppler-Free LMR via Saturation Spectroscopy experiment.

Let's start with panel A, the Experimental Setup. It's a schematic showing a laser beam being split. The strong pump beam goes through a sample cell, where a magnetic field B B is applied. The weak probe beam travels in the opposite direction through the same cell and is measured by a detector.

Panel B shows the Energy Levels and Transitions. On the left, at B = 0 $B = 0$, we have the unsplit upper and lower energy levels. On the right, at B > 0 $B > 0$, these levels split due to the Zeeman effect into m J ′ m_J' and

m J ″ m_J'' sublevels. The allowed transitions, with Δ m J = 0 $\Delta m_J = 0$ and ± 1 ± 1, are indicated.

Panel C illustrates the concepts of Hole Burning and the Lamb Dip. The top graph shows the population of molecules versus their velocity component, v z v_z. The pump beam "burns a hole" in this distribution at v z = 0 $v_z = 0$. The bottom graph shows the absorption profile as you would tune the laser frequency. You see the broad Doppler profile, but right at the center, at frequency v 0 ν_0, there's a sharp dip—this is the Lamb dip, resulting from the probe beam interacting with the saturated velocity group.

Panel D shows the Resulting Doppler-Free LMR Spectrum. Here, the horizontal axis is the magnetic field B B. Instead of a single broad peak, we now see three incredibly sharp, well-resolved Zeeman components, corresponding to Δ m J = − 1 , 0 , + 1 $\Delta m_J = -1, 0, +1$. The crucial annotation here is "Linewidth much, much less than Doppler Width." This is the ultimate achievement of this technique: resolving the intrinsic structure of the transition, hidden within the Doppler profile.

Page 43:

Alright, let's now shift our focus to the second major technique of this lecture: the Stark Effect, which involves shifting energy levels with electric fields. This is the electrical analogue of the Zeeman effect.

The fundamental requirement is that we must be studying polar molecules. A polar molecule is one that possesses a permanent electric dipole moment, which we denote with the vector d d.

When we place such a molecule in a static electric field, which we'll define as E = E z ^ $\mathbf{E} = E\,\hat{\mathbf{z}}$, there is an interaction energy, U U, given by the dot product: U = − d ⋅ E $U = -\mathbf{d} \cdot \mathbf{E}$.

This interaction causes a shift in the energy of the rotational states, known as the Stark shift. For many molecules, particularly symmetric tops, there is a first-order Stark shift. For a rigid rotor state with J ≥ 1 $J \geq 1$, this shift, Δ E 1 ΔE^1, is given by the equation: Δ E 1 = − d E M J (J + 1)

$$\Delta E^1 = -\frac{d\,E\,M}{J(J+1)}$$

Let's break this down. The shift is proportional to the dipole moment d d, the electric field strength E E, and the magnetic quantum number M M (which here represents the projection of J J on the E-field axis). It is inversely proportional to J J times J + 1 $J + 1$.

Page 44

The first-order Stark effect is not always present. For example, in linear molecules, the first-order shift averages to zero. In such cases, or more generally, we must consider the second-order shift, which is derived from second-order perturbation theory.

The second-order energy shift, Δ E (2) $\Delta E^{(2)}$, is given by the equation:

Δ E (2) = − d 2 E 2 2 B r o t f (J , M) .

$$\Delta E^{(2)} = -\frac{d^2E^2}{2B_{\text{rot}}}\,f(J,M).$$

Let's look at the terms here. The shift is now proportional to the square of the dipole moment, d 2 d^2, and the square of the electric field, E 2 E^2. It is inversely proportional to the rotational constant of the molecule, B r o t B_{rot}. I must emphasize that this B r o t B_{rot} is the rotational constant, typically in units of frequency or energy, and should absolutely not be confused with the magnetic field B B! They are completely different physical quantities.

The term f (J , M) $f(J,M)$ is a dimensionless factor that depends on the specific quantum numbers of the state, and its form is derived from perturbation theory.

The key takeaway is that energy levels can shift either linearly with the E-field (first-order) or quadratically (second-order).

Page 45:

So, what is the practical result of these Stark shifts? What kind of tuning range can we achieve?

As a typical result, with electric fields E on the order of 100 kilovolts per centimeter—which are very strong fields but achievable in the lab—the transition frequencies of molecules can be swept, or tuned, by tens to hundreds of megahertz.

While this tuning range is generally smaller than the multi-gigahertz or even terahertz-equivalent shifts seen in LMR, a range of hundreds of megahertz is still more than enough to bridge the gap between a fixed laser line and a

nearby molecular transition, making Laser Stark Spectroscopy a very powerful and practical high-resolution technique.

Page 46:

This diagram illustrates the Stark effect on the energy levels of a J = 1 $J = 1$ state. The plot shows Energy, E E, on the vertical axis versus the Electric Field Strength, E E, on the horizontal axis.

At zero electric field on the left, the three M M sublevels (M = 0 , + 1 , − 1 $M = 0, +1, -1$) are degenerate. As the electric field increases to the right, the levels split.

This diagram is interesting because it depicts both linear and quadratic effects. The levels labeled M = + 1 $M = +1$ and M = − 1 $M = -1$ are shown splitting linearly with the field. This corresponds to the first-order Stark effect, as indicated by the annotation "Linear Stark Effect (Δ E ∝ E $\Delta E \propto E$)". This is what you would expect for a symmetric top molecule.

In contrast, the level labeled M = 0 $M = 0$ is shown shifting downwards with a curved trajectory. This represents a quadratic Stark effect, where the energy shift is proportional to E 2 E^2. This is the behavior expected for a linear molecule.

So this diagram combines the two types of behavior you might encounter. In any given molecule, you will typically be dominated by one or the other. The key point is that the electric field breaks the degeneracy and allows us to tune the energy levels.

Page 47:

Let's now consider the practical arrangement for performing Laser Stark Spectroscopy.

The heart of the experiment is the electrode geometry used to create the strong, uniform electric field. This typically consists of two highly polished, very flat parallel metal plates, which are separated by a very small and precise distance, d d. This distance is often on the order of just one millimeter. Precision-machined spacers, often made of quartz or sapphire, are used to maintain this gap.

A high voltage, V V, of up to 10 10 kilovolts, is applied across these plates. This creates a very strong electric field, E = V d $E = \frac{V}{d}$. For 10 10 kilovolts over 1 1 millimeter, this gives a field of 10 6 10^6 volts per meter, or 100 100 kilovolts per centimeter, just as mentioned before.

Unlike LMR, the Stark cell is almost always placed *outside* the laser resonator. The reason is that the very small, one-millimeter aperture between the plates would introduce significant diffraction losses if placed inside the laser cavity, which would likely prevent the laser from operating at all.

So, how do we detect the signal?

Page 48:

The detection strategy for Stark spectroscopy is very similar to what we saw for LMR.

First, you sweep the DC electric field slowly, in a quasi-static manner. Second, to dramatically improve sensitivity, you superimpose a small AC modulation on top of the DC voltage. You then use a lock-in amplifier, synchronized to this modulation frequency, to detect the signal. Just as in LMR, this yields a signal that looks like the derivative of the absorption line and provides excellent noise rejection.

The sensitivity achieved with this laser-based method is comparable to that of traditional microwave Stark spectrometers, but with the enormous advantage that it can be applied at infrared and far-infrared frequencies, where many more molecules have fundamental transitions.

The measurement accuracy can be extremely high. The electric field can be known very precisely because it depends on the applied voltage, which is easy to measure accurately, and the plate separation, which can be manufactured to high precision. The field is typically known to about one part in ten thousand, or $10 - 4\ 10^{-4}$.

Page 49:

This high accuracy in the field measurement can be leveraged to achieve extraordinary precision in the final result.

By combining the measured resonance condition (the voltage at which the signal appears) with an independent, highly accurate measurement of the fixed laser line frequency, we can determine the absolute frequency of the zero-field molecular transition with incredible precision.

The laser frequency itself can be measured using a high-precision wavemeter, sometimes called a lambda-meter, or through heterodyne techniques where the laser is mixed with a frequency standard.

Using these methods, it's possible to determine the absolute transition frequencies to within 20 to 40 kilohertz. For an infrared transition at, say, 30 terahertz, that's an accuracy of about one part in a billion. This is truly high-precision spectroscopy.

Page 50:

Here is a cross-sectional diagram of a practical Stark spectroscopy cell, illustrating the key components.

We see a large vacuum chamber, the Stark Cell Body, with an IR laser beam passing through it from left to right via two IR windows.

Inside the chamber, at the center, are the two parallel electrodes. We have the top electrode held at a positive high voltage, + V $+V$, and the bottom electrode held at ground. The small, precise separation, d d, is indicated, and the resulting electric field, E = V / d $E = V/d$, is shown by the arrows pointing downwards.

A very important feature shown here are the Guard Rings. These are additional electrodes that surround the main interaction region. Their purpose is to shape the electric field and minimize the non-uniform “fringe fields” that would otherwise exist at the edges of the plates. This ensures that the molecules in the laser beam path experience a highly uniform, well-defined electric field, which is essential for accurate measurements.

This entire assembly is a precision-engineered device designed to apply a very well-known electric field to a gas sample.

Page 51:

Let's look at a case study: the Δ M = 0 $\Delta M = 0$ Stark spectrum of the molecule 14 N H 2 D $^{14}\mathrm{NH_2D}$, which is ammonia with one hydrogen atom replaced by deuterium.

The experiment was performed in the mid-infrared spectral window, between 950 950 and 955 c m − 1 $955\ \mathrm{cm^{-1}}$.

The strategy here was to use multiple different, fixed-frequency laser lines from a C O 2 $\mathrm{CO_2}$ laser. A C O 2 $\mathrm{CO_2}$ laser can be made to lase on many different rotational lines, and each one provides a fixed frequency "anchor point." By using several of these lines, the researchers could interrogate different regions of the molecule's Stark spectrum.

The observations are exactly what we would expect from our conceptual model. Each time a sloped Stark-tuned energy level of the molecule crosses the horizontal line representing a fixed laser frequency, a resonance peak is observed in the spectrum.

From the analysis of the fitted slopes of these tuning transitions, one can extract fundamental molecular properties.

Page 52:

The analysis of the slopes of the Stark-tuning transitions, which is the change in frequency with respect to the electric field, yields two key pieces of information.

First, and most importantly, it allows for the determination of the magnitude of the molecule's permanent electric dipole moment, the absolute value of d d. In this particular study of NH2D, this was done with an uncertainty of ≤ 1 % $\leq 1\%$, which is a very precise measurement of this fundamental molecular property.

Second, a detailed analysis of the Stark spectrum can also be used to extract or refine other important molecular constants, such as the rotational constant, B r o t B_{rot}, and even higher-order terms like centrifugal distortion constants, which account for the fact that a real molecule is not a perfectly rigid rotor.

Page 53: Energy-Field Diagram for a Stark Transition

This diagram shows the actual data from the case study: an "Energy-Field Diagram for a Stark Transition" in 14 N H 2 D $^{14}\mathrm{NH_2\,D}$.

The plot shows Energy or Frequency in units of wavenumbers, or c m − 1 cm^{-1}, on the vertical axis, versus the Electric Field, F F, on the horizontal axis.

The solid blue lines represent the energies of different Stark components of the molecular transition as they tune with the electric field. You can see several lines, some tuning up and some tuning down, with different slopes.

These are labeled with their M quantum numbers, for example, | M | = 4 $|M| = 4$ and | M | = 3 $|M| = 3$.

The two dashed red lines are horizontal, representing two different fixed-frequency lines from the C O 2 CO_2 laser. The experiment consists of finding where the blue lines intersect the red lines. Two such resonance points are marked with red dots.

The callout box gets to the heart of the result. It says: "From fitted slopes (d v / d E dv/dE) of multiple resonances: Determined dipole moment | d | = 1.471 $|d| = 1.471$ Debye, with an uncertainty of less than or equal to 1%." This is the final, high-precision result extracted from this beautiful data.

Page 54:

The techniques of Laser Stark Spectroscopy are powerful, but they can be extended and improved even further. This slide looks at extensions toward broader tuning and even higher, sub-Doppler resolution.

One significant development is the generation of broadly tunable far-infrared radiation by a technique called difference-frequency mixing. Here's how it works: You take the beam from a fixed-frequency CO_2 laser and combine it with the beam from a *tunable* CO_2 waveguide laser. The two beams are focused together onto a very special type of nonlinear device called a metal–insulator–metal, or MIM, diode. This diode acts as a mixer, generating a new frequency which is the difference between the two input laser frequencies.

The result is the production of coherent radiation in the far-infrared, from about 0.1 terahertz up to 3 terahertz, with a narrow linewidth. Crucially, because one of the input lasers is tunable, this output FIR radiation is also tunable. This gives us a tunable source to use in conjunction with Stark spectroscopy.

Page 55: A great example of using such a tunable source is the measurement of the Stark spectrum of the methanol isotope $^{13}CH_3$ OH across a very wide range. Methanol has a very complex spectrum due to the internal rotation, or torsion, of its CH3 group. Having a broadly tunable source allowed the researchers to record many transitions and perform a "global fit" of the entire torsion-

rotation Hamiltonian, leading to a much more complete understanding of the molecule's dynamics.

Another major extension is the implementation of Stark spectroscopy in a molecular beam. Here, a pulsed supersonic jet is used to produce a beam of molecules that is internally very cold—the rotational temperature can be just a few Kelvin. This simplifies the spectrum enormously by collapsing the population into just a few rotational levels.

When you combine this cold molecular beam with a tunable laser and a DC electric field, you can achieve sub-Doppler resolution. This is because in a well-collimated molecular beam, the velocity distribution transverse to the beam direction is very narrow, effectively eliminating Doppler broadening for a laser that crosses the beam at a right angle.

Page 56:

These advanced techniques open up new scientific possibilities. For example, they allow for the determination of dipole moments in vibrationally excited states. These states are often hard to populate and study with traditional microwave spectroscopy, but laser techniques can access them easily.

Looking to the future, there are even more developments on the horizon. One area is UV Stark spectroscopy, using frequency-doubled dye lasers to

perform these experiments on electronic transitions in the ultraviolet. Another area of development is the use of electro-optic crystals, like lithium niobate, for generating and modulating the electric fields, which can offer advantages over the parallel-plate capacitor design in some applications.

Page 57:

This final diagram shows a state-of-the-art experiment that combines several of the advanced techniques we've just discussed: "Sub-Doppler Stark Spectroscopy with Difference-Frequency Generation in a Molecular Beam."

Let's break down this complex setup.

In the top left, we have the tunable FIR radiation source. A fixed-frequency CO_2 laser and a tunable CO_2 waveguide laser are combined and focused onto a MIM diode, producing tunable FIR radiation from 0.1 to 3 THz.

This tunable radiation is then directed into a vacuum chamber.

Inside the chamber, a pulsed supersonic jet creates a molecular beam, which is then collimated by a skimmer to produce a rotationally cold beam of molecules, with a temperature of just a few Kelvin.

The laser beam intersects this molecular beam at a right angle, inside a set of Stark plates where a high-voltage DC electric field, E E, is applied.

The signal is detected, processed, and recorded. A timing diagram at the bottom shows that the laser pulse must be synchronized to fire when the puff of gas from the jet pulse is in the interaction region.

This entire sophisticated apparatus brings together tunable FIR generation, molecular beam cooling, and Stark spectroscopy to achieve the highest levels of resolution and sensitivity.

Page 58:

We have covered a lot of ground today, looking at both Laser Magnetic Resonance and Laser Stark Spectroscopy. To wrap up, let's create a comparative summary of these two powerful techniques.

First, let's consider their common strengths. Both LMR and Stark spectroscopy are built around the strategy of exploiting strong, stable, fixed-frequency lasers, which are readily available in important spectral regions like the mid- and far-infrared.

Both techniques provide exceptionally high sensitivity, which makes them ideal for studying difficult-to-detect species, such as short-lived radicals and other transient molecules found in plasmas, flames, and interstellar space.

And critically, both techniques allow for the precise measurement of fundamental molecular constants. LMR is a prime method for determining g-factors and magnetic hyperfine constants, while Stark spectroscopy is one of the best ways to measure electric dipole moments.

Now, what are the key differences between them?

Page 59:

The differences between LMR and Stark spectroscopy define their distinct areas of application. Let's list the three main ones.

1. Interaction Type: This is the most fundamental difference. LMR is based on a magnetic interaction (the Zeeman effect), while Stark spectroscopy is based on an electric interaction (the Stark effect).

2. Sample Requirements: This follows directly from the interaction type. To have a significant Zeeman effect, a molecule must have a magnetic moment. This generally means it must be paramagnetic—that is, it must have one or more unpaired electrons. For the Stark effect, the molecule must have a permanent electric dipole moment, meaning it must be a polar molecule.

3. Experimental Constraints: The hardware required is quite different. LMR requires a large, powerful, and often superconducting magnet to create a strong, homogeneous magnetic field over the sample volume. Stark spectroscopy, on the other hand, requires a precision-engineered set of high-voltage electrodes with a very small, uniform gap.

This leads to their complementary use in the laboratory.

LMR is the ideal technique for studying open-shell radicals that possess unpaired electrons. Classic examples that have been extensively studied by LMR include CH, OH, and NO.

Page 60:

On the other hand, Stark spectroscopy is the favored technique for studying closed-shell polar molecules, which have no unpaired electrons and thus have a very weak magnetic interaction. It is particularly valuable for molecules that, for whatever reason, lack accessible transitions in the microwave domain, which is the traditional realm of Stark spectroscopy.

Laser Stark spectroscopy opens up the IR and FIR regions for these species.

Finally, let's consider the outlook for these techniques. The future is incredibly bright. The next frontier is to combine these methods with the latest advances in laser technology. Specifically, combining them with frequency-comb-referenced lasers will push the accuracy to the sub-kilohertz level. And combining them with cavity-enhanced techniques, like cavity ring-down spectroscopy, promises to push detection limits down to the parts-per-trillion level. These developments will ensure that LMR and Stark spectroscopy remain at the forefront of high-resolution molecular physics for years to come.

Page 61:

This table provides a final, side-by-side comparative summary of Laser Magnetic Resonance and Laser Stark Spectroscopy. It’s an excellent way to consolidate everything we’ve learned. Let’s go through it row by row.

First, Parameter: Fundamental Interaction. For Laser Magnetic Resonance (LMR), it’s the Zeeman Effect, the interaction with a magnetic field. For Laser Stark Spectroscopy, it’s the Stark Effect, the interaction with an electric field.

Parameter: Applied Field. LMR uses a magnetic field, B B, typically in the range of 0 0 to 2 2 Tesla. Stark spectroscopy uses an electric field, E E, typically 0 0 to 100 100 kilovolts per centimeter.

Parameter: Effective Tuning Range. LMR generally offers a wider tuning range, on the order of 0.1 0.1 to 5 5 gigahertz. Stark spectroscopy offers a smaller range, typically from 10 10 megahertz to about 1 1 gigahertz.

Parameter: Sample Requirement. LMR requires a paramagnetic sample, one with unpaired electrons. Examples are OH, CH, and NO2. Stark spectroscopy requires a polar sample, one with a permanent dipole moment. Examples are methyl fluoride (CH3F), formaldehyde (H2CO), and ammonia (NH3).

Parameter: Primary Measured Property. LMR is used to measure magnetic properties: g g-factors and magnetic hyperfine constants. Stark spectroscopy is used to measure the electric dipole moment, μ μ.

Parameter: Typical Lasers. LMR often uses fixed-frequency far-infrared or mid-infrared lasers like CO2, H2O, or HCN lasers. Stark spectroscopy typically uses fixed-frequency mid-infrared lasers like CO2 or CO lasers.

Finally, Parameter: Resolution and Accuracy. Traditionally, both techniques achieve resolution in the megahertz to kilohertz range. However, when combined with a frequency comb for absolute frequency referencing, the accuracy of both techniques can be pushed into the kilohertz or even sub-kilohertz domain.

This table beautifully encapsulates the similarities, differences, and complementary nature of these two elegant and powerful spectroscopic methods. That concludes our lecture for today. Thank you.