Vol. 2

Chapter 1.6

Page 1:

Good morning, everyone. Welcome back to Physics 608, Laser Spectroscopy. I'm Distinguished Professor Dr M A Gondal, and today, we embark on a new section, Chapter 1.6, where we will explore a remarkably elegant and powerful technique. Over the next lecture, we're going to dive deep into a method that cleverly solves one of the most persistent problems in the spectroscopy of complex environments like plasmas.

Page 2:

The topic for today, as you see emblazoned on the slide, is Velocity-Modulation Spectroscopy. This technique, often abbreviated as VMS, is a beautiful example of how one can use the fundamental principles of electromagnetism and the Doppler effect, combined with some clever electronic signal processing, to achieve extraordinary chemical selectivity. It allows us to pluck the spectral signature of a single type of particle—an ion—out of a veritable soup of other species. By the end of this lecture, you will understand not only how it works, but also appreciate the ingenuity behind its development. Let's begin by first understanding the problem that VMS was invented to solve.

Page 3:

Alright, let's set the stage by considering the challenges and motivation behind Molecular Discharge Spectroscopy. The context here, as the slide notes, is performing spectroscopy *inside a glow discharge cell*.

Now, why would we want to do this? A glow discharge is a type of plasma, and it's an incredibly useful tool for chemists and physicists. As the first bullet point states, a glow-discharge tube is routinely employed to create molecular ions and reactive radicals. We start with a stable, neutral precursor gas—think of something simple like hydrogen, H two, or carbon monoxide, C O. We then apply a strong electric field across the gas at low pressure, which ignites a plasma. Inside this energetic environment, collisions with high-energy electrons, subsequent reactions between ions and molecules, and fragmentation processes create a rich variety of new chemical species. Many of these, particularly molecular ions, are highly reactive and transient. They are fundamentally important in fields ranging from astrophysics and interstellar chemistry to plasma processing for semiconductor manufacturing.

The challenge, however, is that while we can easily *create* this fascinating chemical zoo, studying its inhabitants is another matter entirely. The resulting plasma, as the second bullet point begins to explain, is a complex mixture. Let's look at what's actually in there.

Page 4:

So, what does this plasma contain? The slide lists four major categories of particles, and it's this mixture that creates our central problem.

First, you have the neutral parent molecules. If we started with carbon monoxide, there's still a large amount of neutral CO floating around.

Second, you have numerous neutral fragments. The energetic electrons can break the parent molecules apart. So in a CO discharge, you might get

neutral carbon and oxygen atoms. These are often highly reactive species called radicals.

Third, you have positive ions, or cations. An electron impact can knock an electron off a parent molecule, creating, for example, the CO-plus cation.

And fourth, you have negative ions, or anions, formed by electron attachment.

Now, imagine we shine a tunable laser through this discharge cell and record an absorption spectrum. What do we see? The next bullet point is the crux of the issue. The absorption, or perhaps emission, lines of all these different chemical species overlap heavily in the spectral domain. This creates a dense, congested forest of spectral lines.

Why does this happen? The slide notes a few key reasons. Many different molecules, especially if they have similar masses or structures, will possess similar rotational constants. This means their rotational energy levels are similarly spaced, and their spectra will lie right on top of each other. Furthermore, you have different vibrational bands of different isotopologues—for example, carbon-12-oxygen-16-plus versus carbon-13-oxygen-16-plus—crowding the same wavenumber region.

Page 5:

The situation is made even worse by another fundamental physical effect: Doppler broadening. At the typical cell temperatures, around 300 Kelvin, the random thermal motion of all these species causes their absorption lines to be broadened. This broadening smears out any fine details and causes closely spaced lines to blend together into unresolved lumps.

So we arrive at the final, crucial point on this slide: A purely optical identification of "who owns which line" becomes extremely difficult, if not impossible. You have a spectrum, a plot of absorption versus frequency, that is a complete jumble. It's the sum of the spectra from the parent neutrals, the fragment neutrals, the positive ions, and the negative ions, all broadened and overlapping. The information about the highly interesting ions we want to study is in there, but it's completely buried under the much stronger signals from the far more abundant neutral species. We need a trick. We need a method that provides an additional layer of selectivity to untangle this mess.

Page 6:

This slide provides an excellent visual summary of the challenge we've just discussed. Let's break it down.

At the top, we see a diagram labeled "Spectroscopy Inside a Glow Discharge Cell." On the left is the cathode, which is held at a negative potential, and on the right is the anode, at a positive potential. This creates an electric field between them. The space in between is filled with the plasma, which is depicted as a veritable zoo of particles. The legend helps us identify them: we have positive ions or cations, shown in blue with a plus sign; negative ions or anions, in red with a minus sign; neutral parent molecules, as gray double-spheres; and neutral fragments, as gray single spheres. All of these particles are mixed together.

Now, at the bottom, we see a conceptual graph of what a conventional absorption spectrum of this plasma might look like. The vertical axis is

Absorption or Intensity, and the horizontal axis is Wavenumber, in units of inverse centimeters. The graph shows several individual absorption spectra—a green one, an orange one, a gray one—each corresponding to one of the species in the plasma. When we measure the total absorption, we don't see these individual spectra. Instead, we see their sum, which is the messy, broad feature outlined by the bracket labeled "Spectral Congestion." You might be looking for the tiny absorption signature of a rare molecular ion, but it's completely swamped by the broad, strong absorption of a neutral molecule. This is the problem we need to solve.

Page 7:

So, this brings us to the clear need for a species-selective technique. The limitations of conventional approaches make it plain that just taking a simple absorption spectrum is not enough.

The first bullet point highlights a key failing of a simple, static DC discharge. A DC, or direct current, discharge does not imprint any kind of unique signature on the ionic absorption lines. The spectrum simply shows all species simultaneously. There is no intrinsic label that says "this line belongs to an ion" and "that line belongs to a neutral." It's a free-for-all.

One might think, "Well, can't we be clever? Can't we change the chemistry?" As the second bullet point notes, we could try chemical substitution—changing the precursor gas—or varying the pressure or temperature. While these methods can sometimes help, they are often a blunt instrument. They might shift the line positions, but only very slightly. An assignment based on a tiny shift in a crowded spectrum is often

ambiguous and unconvincing. We need something more direct, more definitive.

Page 8:

Let's consider another approach. If we want to distinguish ions from neutrals, why not use their charge? We could try to electrically extract the ions from the plasma and send them into a mass spectrometer. This would certainly tell us which ions are present. However, as the first bullet point here explains, this method has a major drawback for *in situ* spectroscopy. The process of extraction fundamentally interrupts the optical observation and, more importantly, it completely destroys the plasma environment you are trying to study. You're no longer probing the pristine state of the plasma; you're just analyzing its remnants. This is not true *in situ* spectroscopy.

Therefore, we arrive at the conclusion articulated in the second bullet point, which is really the design goal for the technique we're about to learn. What is highly desirable is an *optical method* that can effectively "tag" the ionic absorbers while leaving the neutral absorbers completely untagged. And it must do this without interrupting or disturbing the discharge itself. How on earth can we achieve this?

Page 9:

Here is the solution, in preview. The answer is Velocity-Modulation Spectroscopy, or VMS. This technique is brilliantly simple in its conception.

The first bullet point lays out the core mechanism. VMS introduces a *periodic drift velocity* that is applied *exclusively* to the charged particles. How does it do this? By rapidly switching the polarity of the discharge voltage. Imagine the anode and cathode swapping their positive and negative identities back and forth, thousands of times per second. We'll see that this sets the ions, and only the ions, into a periodic, oscillating motion.

The second bullet point describes the spectroscopic consequence, which is the key to the whole technique. Because only the ionic absorbers are moving back and forth with this periodic drift velocity, only they experience a *first-order Doppler modulation*. Their absorption frequencies are periodically shifted up and down. The neutral species, on the other hand, are unaffected by the switching electric field. They just continue their random thermal motion and remain at thermal equilibrium, with a zero average drift velocity. Their absorption frequencies do not get modulated.

This modulation is the "tag" we were looking for. It's a unique signature imprinted only on the ions, allowing us to distinguish them from the sea of neutral molecules.

Page 10:

So, let's build the physics from the ground up. We'll start by considering the electric field-driven ion drift in a standard glow discharge. How do we establish a macroscopic ion velocity?

It's quite simple, as the first point describes. We apply an external voltage, which we'll call V sub ext, between the anode and the cathode. Let's say these electrodes are separated by a distance d

.

As you know from introductory electromagnetism, this applied voltage creates an electric field in the region between the electrodes. To a first approximation, the magnitude of this resulting electric field, capital E, is given by the simple relation shown in the equation:

E = V ext d

$$E = \frac{V_{\text{ext}}}{d}$$

E equals V sub ext over d.

E = V ext d

$$E = \frac{V_{\text{ext}}}{d}$$

This electric field permeates the plasma and will exert a force on any charged particles within it.

Page 11:

Now, what does this electric field, capital E, which has units of Volts per meter, do? It accelerates charged particles.

The force, capital F, on an ion with charge q q is given by the fundamental Lorentz force equation, simplified for the electric field:

F = q E

$$F = qE$$

The charge q q is, of course, quantized. As the next bullet point states, for a singly charged positive ion, a cation, the charge q q is equal to + e $+e$, where e e is the elementary charge. For a singly charged negative ion, an anion, q q is equal to − e $-e$.

Now, a crucial point. An ion in the discharge is not in a vacuum. It's constantly bumping into the far more numerous neutral buffer gas atoms. So, it doesn't accelerate indefinitely. After being accelerated by the field, it quickly collides with a neutral, loses some of its directed momentum, and then gets accelerated again. This process repeats over and over.

The slide notes that the mean free time between these collisions, τ c o l l τ_{coll}, is incredibly short, on the order of 10 − 12 10^{-12} seconds, or one picosecond. The result of this rapid acceleration and collision cycle is that the ion reaches a steady-state *average drift velocity*, which we denote as v D v_D. This drift velocity is a net, directed motion superimposed on the ion's random thermal motion.

Page 12:

This brings us to a key relationship for the average drift velocity, v D v_D. In the regime of low electric fields, which is typically the case in a glow discharge, this drift velocity is directly proportional to the electric field magnitude, E E. The equation is:

v D = μ ⋅ E

$$v_D = \mu \cdot E$$

Let's break this down. v D v_D is the drift velocity in meters per second. E E is the electric field in Volts per meter. The constant of proportionality, the Greek letter μ μ, is a critically important parameter called the *ionic mobility*.

As the first bullet point explains, the ionic mobility, μ μ, has units of meters squared per Volt per second. It's a measure of how easily an ion can move through a specific buffer gas under the influence of an electric field. Its value depends on the properties of the ion—its mass and charge—and on its collision cross-section with the buffer gas atoms. A larger, heavier ion will generally have a lower mobility than a smaller, lighter one.

Now for the most important contrast, stated in the final bullet point. What about neutral molecules? They have no net electric charge; q = 0 $q = 0$. Therefore, they experience no net electric force. And if there is no force, there can be no drift velocity. So, on average, v D v_D is zero for all neutral species. This is the fundamental physical difference that Velocity-Modulation Spectroscopy exploits.

Page 13:

This diagram provides a perfect visual representation of the concept of electric field-driven ion drift that we've just been discussing.

Let's orient ourselves. On the left, we have the anode, marked with a plus sign, indicating it's at a higher potential. On the right, we have the cathode, marked with a minus sign, at a lower potential. This setup creates an electric field, represented by the vector E and the dashed orange lines, pointing from the anode to the cathode, so from left to right.

Now look at the particles within this field, which are identified by the particle legend.

First, consider a positive ion, or cation, shown as a blue circle with a plus. Its charge q q is positive. The force on it, F = q E $F = qE$, is in the same direction as the electric field. You can see the red arrow labeled F F pointing to the right. As a result, this cation acquires a drift velocity, v D v_D, also pointing to the right, toward the cathode.

Next, look at the negative ion, or anion, shown as a red circle with a minus. Its charge q q is negative. The force on it, F = q E $F = qE$, is in the *opposite* direction to the electric field. The force vector F F points to the left. Consequently, the anion acquires a drift velocity v D v_D to the left, toward the anode.

Finally, and most importantly, look at the neutral molecule, shown as a gray circle. Its charge is zero. It feels no electric force. Therefore, its average drift velocity, v D v_D, is zero. It just sits there, undergoing random thermal motion, but with no net directional movement.

This simple picture contains the entire physical basis for the selectivity of VMS. The electric field sorts the particles by charge, imparting a directional velocity to ions but not to neutrals.

Page 14:

Now that we've established that ions have a drift velocity, we need to connect this motion to an observable spectroscopic effect. This brings us back to the Doppler effect. Let's revisit it, specifically for the case of moving

absorbers, using a straightforward derivation based on the principle of phase invariance.

As the first bullet point states, consider a monochromatic plane laser wave propagating along the x-axis of our discharge cell. As we know, such a wave can be described by its space-time phase, capital Phi. The equation for the phase is:

Φ = k x − ω t

$$\Phi = kx - \omega t$$

Let's define these terms to be perfectly clear.

Capital Phi is the phase of the wave. k

is the magnitude of the wave-vector, which is equal to two pi divided by the wavelength, lambda. It tells us how the phase changes with position and has units of radians per meter. x

is the position along the axis. Omega, which is equal to two pi times the frequency nu, is the angular frequency of the laser light in the laboratory frame. It has units of radians per second. And t

is time.

This phase is what a stationary observer at position x would measure. But what about our moving ion?

Page 15:

Here is the critical step in the derivation. For an atom or ion to absorb a photon from the light wave, the process must be resonant. This means that

from the ion's perspective, as it moves, the phase of the electromagnetic wave it experiences must remain stationary at the moment of absorption. In other words, the total time derivative of the phase, as seen by the moving ion, must be zero.

The slide expresses this condition as:

d Φ d t = 0

$$\frac{d\Phi}{dt} = 0$$

Let's perform this differentiation. The phase is Φ = k x − ω t $\Phi = kx - \omega t$. Remember that for the moving ion, its position x x is also a function of time. Using the chain rule, the total derivative is d Φ d t = k d x d t − ω $\frac{d\Phi}{dt} = k\frac{dx}{dt} - \omega$. What is d x d t $\frac{dx}{dt}$? It's simply the ion's velocity along the x-axis, which is its drift velocity, v D v_D.

So the condition becomes:

k v D − ω = 0

$$kv_D - \omega = 0$$

Now, we need to be very careful here. The frequency ω ω in this equation is the frequency in the *ion's rest frame* that it absorbs. This is the natural transition frequency of the ion, which we will call omega-naught, ω 0 ω_0. The frequency ω ω in the original phase equation Φ = k x − ω t $\Phi = kx - \omega t$ was the laser frequency in the *lab frame*. A more rigorous derivation equates the lab-frame frequency Doppler-shifted into the ion's frame with

the ion's rest-frame transition frequency. This leads to the result shown in the next step.

Let's solve for the laboratory frequency, omega, that is absorbed by the drifting ion. The result is:

ω = k v D + ω 0

$$\omega = k v_{\mathrm{D}} + \omega_0$$

Here, as the last point clarifies, ω 0 ω_0 is the natural, unshifted transition frequency in the ion's own rest frame. This equation tells us that an ion moving with drift velocity v D v_{D} will not absorb light at its rest frequency ω 0 ω_0, but rather at a lab frequency ω ω that is shifted by an amount k v D $k v_{\mathrm{D}}$.

Page 16:

From the equation on the previous slide, ω = k · v D + ω 0 $\omega = k \cdot v_{\mathrm{D}} + \omega_0$, we can directly define the first-order Doppler shift. The shift, which we'll call Delta omega, is the difference between the observed lab-frame frequency, omega, and the ion's rest-frame frequency, omega-naught. As the equation shows:

Δ ω = ω − ω 0 = k · v D

$$\Delta\omega = \omega - \omega_0 = k \cdot v_{\mathrm{D}}$$

This is the fundamental relationship. The drift velocity v D v_{D} directly causes a frequency shift Δ ω $\Delta\omega$.

Now let's consider the sign convention, which is very important. Let's assume our laser beam is propagating in the positive x-direction.

If an ion has a positive drift velocity v D v_D, meaning it is moving *toward* the laser source (or, more accurately, moving counter-propagating to the laser beam), the term k · v D $k \cdot v_D$ is positive. The absorption frequency ω ω is higher than ω 0 ω_0. This is a blue-shift.

Conversely, if an ion has a negative drift velocity v D v_D, meaning it is moving *away* from the laser source (co-propagating with the beam), the term k · v D $k \cdot v_D$ is negative. The absorption frequency ω ω is lower than ω 0 ω_0. This is a red-shift.

So, the direction of the ion's drift relative to the laser beam's propagation determines whether the absorption line is blue-shifted or red-shifted. And for neutrals, since v D v_D is zero, Δ ω $\Delta\omega$ is zero. There is no shift.

Page 17:

Alright, we have the theory. The ion's drift velocity causes a Doppler shift. But is this effect large enough to be useful? Let's plug in some typical numbers from a laboratory discharge to get a quantitative feel for the magnitude of this shift.

First, let's consider the laser. A common wavelength for studying molecular vibrations is in the near-infrared. The slide gives an example of λ = 1.55 μ m $\lambda = 1.55\,\mu$m. From this, we can calculate the magnitude of the wave-vector, k k.

k = 2 π λ

$$k = \frac{2\pi}{\lambda}$$

Plugging in the numbers, k = 2 π 1.55 × 10 − 6 m $k = \frac{2\pi}{1.55\times10^{-6}\,\text{m}}$. This gives a value of approximately 4.05 × 10 6 m − 1 $4.05 \times 10^{6}\,\text{m}^{-1}$.

Next, we need the ionic mobility, μ μ. The slide provides a typical value for light molecular ions, like H 3 + H_3^+, in a helium buffer gas. A reasonable estimate for μ μ is:

μ ≈ 1.5 × 10 − 3 m 2 V ⋅ s

$$\mu \approx 1.5 \times 10^{-3}\ \frac{\text{m}^2}{\text{V}\cdot\text{s}}$$

Now we have k k and μ μ. The last piece of the puzzle is the electric field, E E.

Page 18:

Let's continue our calculation with a realistic electric field. A typical discharge might have a voltage of 300 Volts applied across a distance of 10 centimeters, or 0.1 meters. So, the electric field E is:

E = 300 Volts 0.1 meters = 3.0 × 10 3 Volts/m

$$E = \frac{300\ \text{Volts}}{0.1\ \text{meters}} = 3.0 \times 10^{3}\ \text{Volts/m}$$

Now we can calculate the drift velocity, v D v_D, using the relation v D = μ E $v_D = \mu E$.

v sub D is approximately one point five times ten to the minus three, times three times ten to the three, which gives a drift velocity of about four point five meters per second. This is a very modest speed, much slower than the thermal velocities of the ions.

With this drift velocity, we can finally calculate the resulting angular-frequency shift, Delta omega, using our Doppler shift formula, Δ ω = k v D $\Delta\omega = k\, v_D$.

Delta omega is approximately four point zero five times ten to the six inverse meters, times four point five meters per second. This gives an angular frequency shift of about one point eight times ten to the seven radians per second.

To make this more intuitive, let's convert it to a frequency shift in Hertz, which we call Delta nu. We just divide Delta omega by two pi.

Δ v = Δ ω 2 π ≈ 2.9 MHz

$$\Delta\nu = \frac{\Delta\omega}{2\pi} \approx 2.9 \text{ MHz}$$

Now, here is the crucial comparison. The last bullet point notes that the thermal Doppler *width*—the broadening due to random thermal motion—for the same ion at 300 Kelvin is typically around 200 Megahertz. Our drift-induced *shift* is about 3 MHz. So, the drift shift is only about one to two percent of the total linewidth.

This might seem small, but it's actually perfect.

Page 19:

This small but significant shift is exactly what we need. It is large enough to modulate the absorption signal in a detectable way, but it is not so large that it moves the absorption frequency completely outside of the line shape envelope.

Think of it this way: the absorption line is like a broad hill, about 200 MHz wide. Our velocity-modulation technique makes this entire hill jiggle back and forth by about 3 MHz. If you're sitting on the side of the hill with your fixed-frequency laser, you will see your absorption signal go up and down as the hill moves under you. This modulation is what we will detect. If the shift were much larger than the linewidth, the line would jump completely past the laser frequency, which would lead to a different type of signal. The fact that the shift is a small fraction of the width is ideal for the lock-in detection scheme we are about to discuss.

Page 20:

So far, we've considered a static DC field, which causes a static Doppler shift. But a static shift doesn't help us distinguish anything. The key, as we previewed, is modulation. This brings us to the core concept of Velocity Modulation, which involves introducing an alternating, or AC, electric field.

As the first bullet point states, we simply replace the DC voltage source with an AC voltage source. The external voltage, V sub ext, now becomes a function of time, t

. A simple and common choice is a sinusoidal voltage:

$$V_{ext}(t) = V_0 \sin(2\pi f t)$$

$$V_{ext}(t) = V_0 \sin(2\pi f t)$$

Let's look at the terms here. V $_0$

is the amplitude of the voltage, and the slide gives a typical value of around 300 Volts. f

is the modulation frequency. This is a crucial parameter, and typical values are in the audio frequency range, from 1 kilohertz up to about 50 kilohertz.

Naturally, if the voltage is oscillating, so is the electric field. The electric field E also becomes a function of time:

E (t) = V 0 sin (2 π f t) d

$$E(t) = \frac{V_0 \sin(2\pi f t)}{d}$$

This time-varying electric field will now drive a time-varying force on the ions.

Page 21:

What is the effect of this oscillating electric field on the ions? The ionic drift velocity, v D v_D, will now follow the field.

The first bullet point notes a very important condition: the response is quasi-instantaneous. This is because the mean time between collisions, which we called τ i o n τ_{ion} or τ c o l l τ_{coll}, is on the order of picoseconds. The period of the AC field, which is one over f f, is on the order of microseconds. Since τ i o n τ_{ion} is much, much less than 1 f $\frac{1}{f}$, the ion population can respond to the changing field essentially instantly.

So, the drift velocity v D v_D also becomes a sinusoidal function of time:

v D (t) = μ E (t) = v 0 sin ⁡ (2 π f t)

$$v_D(t) = \mu E(t) = v_0 \sin(2\pi f t)$$

Here, v 0 v_0 is the *amplitude* of the drift velocity, given by μ μ times the amplitude of the electric field, which is μ μ times V 0 V_0 over d d.

Now we connect this back to spectroscopy. The absorption frequency experienced by the moving ion, which we found was ω = ω 0 + k v D $\omega = \omega_0 + k v_D$, is therefore also dynamically modulated. It becomes:

ω (t) = ω 0 + k v 0 sin ⁡ (2 π f t)

$$\omega(t) = \omega_0 + k v_0 \sin(2\pi f t)$$

The absorption frequency of the ion is wiggling back and forth around its central rest frequency, ω 0 ω_0, at the modulation frequency f f.

And, as always, we return to the critical point of selectivity, stated in the final bullet. What about the neutral species? Since their charge q q is zero, they do not partake in this modulation at all. Their absorption frequency remains fixed at ω 0 ω_0. This is the key that unlocks the spectrum.

Page 22:

So, let's pause for a moment and recap the central physical principle we've established. By applying an AC electric field to a plasma, we cause the ions to oscillate back and forth, which in turn causes their absorption frequencies to be modulated via the Doppler effect. Neutral species remain completely unaffected. The ionic lines now carry a "tag"—this modulation at

frequency f f—while the neutral lines do not. The next question is, how do we design a detection system to read this tag and ignore everything else?

Page 23:

This brings us to the detection scheme, which relies on a powerful analogy and a standard piece of electronic instrumentation. The title says it all: "Frequency Modulation Analogy & Lock-In Detection." Our goal is to extract the unique ionic signature.

The first bullet point makes a very insightful connection. The periodic shifting of the ion's absorption frequency, ω (t) = ω 0 + Δ ω max sin ⁡ (2 π f t) $\omega(t) = \omega_0 + \Delta\omega_{\text{max}}\sin(2\pi f t)$, is mathematically equivalent to performing small-amplitude frequency modulation, or FM, on a laser beam and probing a stationary absorption line.

Many of you may be familiar with FM spectroscopy. It's a high-sensitivity technique where you modulate the laser's frequency and look for signals at the modulation frequency. Here, we are doing something conceptually similar, but instead of modulating the laser, we are modulating the *molecule's* transition frequency. The mathematical formalism and the detection principle are largely the same.

Page 24:

So how do we put this principle into practice? Let's walk through the process.

First, we take our tunable laser, with its frequency ω L ω_L, and we scan it slowly across the spectral region of interest. At the same time, the discharge is being modulated at the frequency f f. When the laser frequency ω L ω_L is near an *ionic* absorption line, the transmitted laser intensity, I (t) $I(t)$, will pick up the modulation. It will contain signal components, or sidebands, at the modulation frequency f f. Why? Because as the absorption line wiggles back and forth underneath the laser, the amount of light absorbed will vary periodically.

Next, we take this time-varying intensity signal, I (t) $I(t)$, and we feed it into a phase-sensitive amplifier, more commonly known as a lock-in amplifier. This instrument is the heart of the detection scheme. We also provide the lock-in amplifier with a reference signal from the same function generator that is driving the discharge modulation at frequency f f. The lock-in amplifier is an incredibly powerful electronic filter. It is designed to extract *only* the component of the input signal I (t) $I(t)$ that has the same frequency, f f, and a fixed phase relationship with the reference signal. In doing so, it mercilessly rejects all other frequency components—DC offsets, one-over-f noise, random fluctuations, and, most importantly, the signals from neutral species.

Finally, what does the output of the lock-in amplifier look like? As the last bullet states, for a sufficiently small modulation amplitude k · v 0 $k \cdot v_0$ (which we established is the case), the lock-in output is directly proportional to the *first derivative* of the unmodulated ionic absorption line shape. So, instead of a peak, you see a characteristic bipolar, "S"-shaped signal as you scan the laser across the line.

Page 25:

And now for the grand finale of the detection scheme. What happens when the laser scans over a neutral-molecule line?

As we've stressed repeatedly, the neutral lines lack modulation. Their absorption frequencies are static. Therefore, the transmitted intensity I (t) $I(t)$ shows no variation at the modulation frequency f f. When this unmodulated signal is fed into the lock-in amplifier, the lock-in, which is looking exclusively for signals at frequency f f, finds nothing.

The result is that the neutral lines yield a zero signal at the lock-in output. They are rendered completely invisible. They therefore vanish from the detected spectrum.

This is the incredible power of VMS. It acts as a perfect filter, throwing away the overwhelming signal from the neutral species and allowing us to see the much weaker, once-hidden spectrum of the ions alone.

Page 26:

The power of Velocity-Modulation Spectroscopy doesn't even stop there. It has another trick up its sleeve, allowing us to perform Phase Discrimination Between Positive and Negative Ions. This gives us yet another layer of information for free.

The core idea is based on opposite Doppler shifts leading to opposite lock-in phases. Let's think it through, as described in the first bullet point. Let's say our laser beam is propagating from left to right. Under an electric field

that is also pointing from left to right (anode on the left, cathode on the right), what happens?

Positive ions, the cations, drift *with* the field, toward the cathode. If they are moving away from the laser source, they experience a red-shift. Let's adjust our convention for a moment and just consider the sign of the velocity relative to a fixed axis. The ions drift with velocity v D v_D. The Doppler shift is proportional to k · v D $k \cdot v_D$. Let's say this is a positive shift, + k · v D $+k \cdot v_D$.

Now, the negative ions, the anions, drift *against* the field, toward the anode. Their velocity is in the opposite direction, − v D $-v_D$. Their Doppler shift will therefore be − k · v D $-k \cdot v_D$.

The key is that the sign of the Doppler shift is opposite for cations and anions.

This leads directly to the conclusion in the second bullet point. Because the sign of the frequency shift Δ ω $\Delta\omega$ reverses, the phase of the AC signal seen by the detector is inverted. This means the phase of the signal going into the lock-in amplifier differs by π π radians, or 180 degrees, between cations and anions.

Page 27:

What is the practical outcome of this 180-degree phase difference? It provides a simple and direct way to distinguish between positive and negative ions.

As the first bullet point explains, we can display the lock-in output on an oscilloscope or record it with a data acquisition system that uses dual-channel demodulation. A lock-in amplifier can measure the signal component that is "in-phase" with the reference (the X component) and the component that is 90 degrees out of phase, or in "quadrature" (the Y component). By adjusting the reference phase, we can set it up so that one type of ion appears purely in the X channel.

For example, we can adjust the phase so that the first-derivative signals from cation lines appear as positive-going peaks in the final spectrum. Because the anion signals are 180 degrees out of phase, they will consequently appear as negative-going peaks. If the concentrations and transition strengths were identical, the peaks would be mirror images of each other.

The conclusion, stated in the final bullet, is truly remarkable. Velocity-Modulation Spectroscopy not only isolates all ionic species from the neutrals, but it also *immediately* and *unambiguously* distinguishes the charge polarity—positive versus negative—without any additional instrumentation or complicated analysis. It's all contained in the sign of the lock-in signal.

Page 28:

Now that we understand the principles, let's look at the canonical experimental arrangement for VMS. What are the key optical and electrical components you would find in a typical lab setup?

First, the optical components. The heart of any laser spectroscopy experiment is, of course, the laser. As the first bullet point states, we need a narrow-linewidth, tunable laser. Historically, color-center lasers or dye lasers were common. Today, diode lasers are very frequently used. The laser beam traverses the plasma column inside the discharge tube. To increase the signal, which depends on the path length, the beam often passes through the plasma multiple times. This is achieved using a multipass arrangement like a Herriott cell, which uses two spherical mirrors to fold the beam path back and forth through the sample, increasing the effective absorption path length from centimeters to many meters.

Next is the discharge tube itself. As the second bullet describes, a typical tube might be about 50 centimeters long with an inner diameter of about 1 centimeter. A critical feature is the windows at each end. To minimize reflection losses of the laser beam, the tube is fitted with transparent windows mounted at Brewster's angle. This ensures that light of a specific polarization can enter and exit the cell with nearly 100 percent transmission.

Page 29:

Now for the electrical components that drive the whole process.

The most specialized piece of equipment is a high-voltage, polarity-reversal switch. This is the device that takes a low-voltage signal from a function generator and uses it to apply a high-voltage AC waveform to the discharge electrodes. As the first bullet notes, it can supply either a sinusoidal or a

square-wave AC voltage, typically up to frequencies of 50 kilohertz, and must be able to handle currents of several amperes to sustain the plasma.

Next, after the laser passes through the cell, we need to detect it. This requires a fast photodetector. Its electronic bandwidth must be significantly larger than the modulation frequency f f—the slide says $\gg 2$ f $\gg 2f$ to be safe—so that it can accurately follow the intensity variations caused by the ionic modulation. The electrical signal from this detector is the input to our lock-in amplifier, which, as we've said, is referenced to the same function generator driving the discharge.

Finally, we need to record the spectrum. A computer or a chart recorder acquires the DC output signal from the lock-in amplifier. This acquisition is synchronized with the tuning of the laser, which is typically done by applying a voltage ramp to a piezo element or by changing the injection current of a diode laser. The result is a plot of the lock-in signal versus laser frequency.

Page 30:

Here we have a complete block diagram of the canonical experimental arrangement for Velocity-Modulation Spectroscopy. Let's trace the signal flow from start to finish to solidify our understanding.

Start at the bottom left with the **Function Generator**. This is the master clock of the experiment. It produces a stable sine wave at the modulation frequency, f f.

This reference signal splits and goes two ways. First, an orange line shows it going to the **High-Voltage AC Amplifier**. This amplifier boosts the signal

and applies the high-voltage AC waveform to the electrodes of the **AC Discharge Tube**, which contains the plasma. Second, another orange line shows the reference signal going to the R e f

input of the **Lock-in Amplifier**. This tells the lock-in what frequency to "lock on" to.

Now let's follow the optical path. At the top left, the **Tunable Laser** (which could be a diode, dye, or other type) produces a laser beam. The frequency of this laser is slowly scanned, as indicated by the "Laser Frequency Scan" input arrow. The beam, with frequency v L ν_L, enters the discharge tube. The diagram shows it making multiple passes using what are labeled as "Multipass Herriott Cell Mirrors," which increases the interaction length.

After traversing the plasma, the transmitted laser beam exits and strikes the **Fast Photodetector**. The detector's bandwidth must be much greater than 2 f $2f$. The detector converts the optical intensity into a time-varying electrical signal, I (t) $I(t)$, which contains the modulation information from the ions.

This signal, I (t) $I(t)$, is fed into the signal input of the **Lock-in Amplifier**. The lock-in, which is a phase-sensitive detector, compares the input I (t) $I(t)$ with the reference signal f f. It filters out everything except the signal component at frequency f f and outputs a clean, DC voltage, S (v) $S(\nu)$, which is the "Lock-in Output".

Finally, this output signal is sent to the **Data Acquisition & Control Computer**. The computer records the lock-in signal S (v) $S(\nu)$ as a function of the scanned laser frequency v ν. The result, as shown in the

small inset, is a plot displaying the characteristic first-derivative lineshape of the ionic absorption, free from the neutral background.

Page 31:

Let's now delve into some of the practical implementation details and optimization strategies. First, the electrical considerations, which are crucial for achieving high performance at high modulation frequencies.

The first bullet point advises using low-inductance electrodes and coaxial feedthroughs. Why is this important? At high frequencies, parasitic inductance in the wiring and electrodes can limit the speed at which the voltage can be switched. The rise time of a circuit is often limited by L R $\frac{L}{R}$ or R C RC time constants. To achieve a crisp, fast reversal of the discharge polarity, you need to minimize this inductance and capacitance. Coaxial cables help to control the impedance and shield the signals.

The second point discusses the heart of the high-voltage driver: the electronic switching circuits. To achieve the rapid, high-voltage polarity reversal, these circuits often employ stacks of solid-state transistors like MOSFETs, or, for very high power applications, devices called thyratrons. A well-designed circuit can achieve a full ± 300 Volt ± 300 Volt polarity reversal in less than 5 µ s $5\,\mu$s. This fast switching is essential for operating at high modulation frequencies, which helps to move the detection band away from low-frequency one-over-f noise.

Page 32:

Continuing with practical details, the slide lists some typical operating conditions for a VMS experiment. You might run the discharge with a 300 Volt peak-to-peak AC voltage, drawing about 3 amps of root-mean-square current. The gas pressure is typically low, in the range of 0.5 to 1.0 Torr. Under these conditions, the discharge power can be quite significant, around 1 kilowatt, which means the discharge tube often requires cooling.

Now for some optical considerations. A major advantage of using lock-in detection is its noise rejection capability. As the first point states, any laser intensity noise that occurs at frequencies *outside* the narrow detection bandwidth of the lock-in amplifier does not influence the measurement. It's simply filtered out. However, if the discharge itself causes some residual amplitude modulation (RAM) of the laser beam at the modulation frequency, this can create a spurious background signal. To combat this, one can use a balanced detection scheme, which can further suppress this type of noise.

The second optical point reiterates the importance of the multipass arrangement. By increasing the effective path length to several meters, it dramatically enhances the total ionic absorbance. The absorbance, A A, is given by the natural logarithm of the ratio of incident intensity I 0 I_0 to transmitted intensity I I. This enhancement brings very weak absorption signals, on the order of A A being around 10 − 3 10^{-3}, into a readily detectable range.

Page 33:

Ultimately, the performance of any spectroscopic technique is judged by its signal-to-noise ratio, or SNR.

For VMS, when properly optimized, the dominant noise source is often the fundamental shot noise of the photons themselves. In this shot-noise-limited regime, the SNR has a very favorable scaling law. As the first bullet point notes, the SNR improves as the square root of τ i n t τ_{int}, where τ i n t τ_{int} is the lock-in time constant.

The time constant, or integration time, of the lock-in amplifier sets the detection bandwidth. A longer time constant means a narrower bandwidth, which filters out more noise, thus improving the SNR. A typical value for τ i n t τ_{int} is around 100 milliseconds. Doubling the time constant doesn't double the SNR, but it improves it by a factor of 2 $\sqrt{2}$.

How good can the SNR get? The second bullet point gives a very impressive real-world result. Researchers have reported achieving an SNR greater than one thousand for weak absorption lines of the CO-plus cation. And this was achieved with an averaging time of only a few seconds. This demonstrates the exceptional sensitivity of the VMS technique.

Page 34:

Let's ground our discussion in a real-world example with a case study: the detection of the band head of the carbon monoxide cation, CO-plus. This is a classic application that beautifully demonstrates the power of VMS.

The specific transition being monitored is an electronic transition in the CO-plus ion. It's the A doublet Pi one-half state, with vibrational quantum

number v ′ = 1 $v' = 1$, being excited from the X doublet Sigma g plus ground state, with v ″ = 0 $v'' = 0$. You don't need to be an expert in molecular term symbols to understand the key point: this is a specific, known transition that belongs only to the CO-plus ion.

The experiment focuses on the region around the R one two one branch band head. A band head is a place in a molecular spectrum where the rotational lines become extremely close together and turn back on themselves. This extreme rotational congestion makes it a perfect, challenging test case for any spectroscopic technique. High resolution is essential.

What were the observational highlights? The velocity-modulated spectrum displays clear, first-derivative line shapes. Crucially, these lines all show a *clean positive polarity*. As we learned, by setting the lock-in phase appropriately, this positive sign immediately and unambiguously confirms that the absorbing species, CO-plus, is indeed a cation.

Page 35:

Continuing with the results of the CO-plus case study, the next bullet point highlights the most dramatic advantage of the technique. The lock-in trace, which is the final VMS spectrum, is completely free from the underlying absorption of neutral molecules.

In a conventional absorption spectrum of a CO discharge, the weak CO-plus signal would be sitting on top of, and completely obscured by, much stronger absorption features from the neutral parent molecule, CO, and potentially from contaminants like O two. These broad, strong neutral

signals create a large, sloping background that makes it impossible to even see the ion's spectrum, let alone measure it accurately.

With VMS, this entire background is eliminated. The lock-in amplifier simply doesn't respond to the unmodulated neutral signals, resulting in a perfectly flat, zero-signal baseline. This "background-free" nature is what enables the high sensitivity and clean interpretation of the ionic spectrum.

Page 36:

This slide is the payoff. It shows the actual data from the CO-plus case study and the contrast is stunning.

Let's first look at the top panel, labeled "(a) Conventional Absorption." The vertical axis is Absorbance, and the horizontal axis is Wavenumber in inverse centimeters, from 2204 to 2211. What you see is a spectrum that is dominated by a broad, slowly varying feature labeled "Broad neutral absorption." Superimposed on this are some weak, sharp wiggles. One of them is pointed out with an arrow labeled "Weak CO^+ Absorption." You can barely see it. This is a perfect illustration of our problem: the signal of interest is almost completely lost in the background.

Now, look at the bottom panel, labeled "(b) Velocity Modulation." The vertical axis is now the VM Signal from the lock-in amplifier, in arbitrary units. The first thing you notice is that the massive, broad neutral background is gone. The baseline is flat at zero. What remains is a beautiful, clean series of sharp features. These are the individual rotational lines of the CO-plus ion. Each line has the characteristic "first-derivative line shape," with a positive lobe and a negative lobe. The spectrum is so

clean that you can resolve the dense structure near the R-branch band head, which is marked around 2210 inverse centimeters.

Comparing panel (a) and panel (b) tells the entire story. VMS has acted like a magic filter, removing the haystack of neutral absorption and leaving behind only the needle of the ionic spectrum.

Page 37:

Of course, to perform laser spectroscopy, you need a laser. So, let's briefly discuss suitable laser sources for VMS across different parts of the electromagnetic spectrum.

First, the Infrared Region, from roughly 1 to 5 micrometers. This region is vital for studying the fundamental vibrational transitions of molecules. Historically, as the first bullet point mentions, color-center lasers were important tunable sources. These were complex, cryogenic systems, for instance using crystals like Lithium Fluoride with F-two color centers, or Sodium Chloride doped with OH-minus. They provided milliwatt-level power and broad tunability.

More common today are semiconductor diode lasers. Specifically, distributed-feedback, or DFB, diode lasers are excellent sources for this region. As the second bullet highlights, they offer very narrow linewidths, typically less than a Megahertz, which is crucial for high-resolution spectroscopy. Furthermore, their frequency can be easily and rapidly tuned by simply modulating their injection current, which makes them ideal for the fine scanning required in a VMS experiment.

Page 38:

Moving to shorter wavelengths, let's consider the Visible and Near-Ultraviolet regions of the spectrum. Probing molecules here allows us to study their electronic transitions, which are typically much stronger than vibrational transitions.

The workhorse for tunable visible lasers for many decades was the dye laser. As the first bullet point notes, using various organic dyes like Rhodamine 6G or Coumarin 480, these lasers could be tuned across the visible spectrum, from roughly 350 to 700 nanometers. The pioneering work extending VMS into the visible range was done with these systems.

More recently, as the second bullet points out, External-Cavity Diode Lasers, or ECDLs, have become a popular alternative. These systems use a standard laser diode as the gain medium but place it inside a larger optical cavity with a grating or other wavelength-selective element. This allows for narrow-linewidth, widely tunable operation. ECDLs now complement or even replace dye laser systems, offering the significant advantages of being solid-state, more compact, and having much simpler alignment and computer control.

Page 39:

What about even shorter wavelengths? There are many interesting radical ions whose strongest electronic transitions lie in the ultraviolet. This slide points to these short-wavelength prospects.

To access the UV region from 250 to 350 nanometers, we can use the technique of frequency doubling. As the bullet point describes, we can take the output of a powerful, tunable laser in the visible, such as a dye laser or a Titanium-sapphire laser, and pass it through a non-linear optical crystal. This process of second-harmonic generation can efficiently convert the light to half its original wavelength, thus doubling its frequency and pushing it into the UV.

Accessing this spectral window is important because many radical ions have strong electronic transitions there, often designated as A-to-X bands. VMS combined with these UV laser sources provides a powerful tool for studying the electronic structure of these highly reactive species.

Page 40:

Velocity-Modulation Spectroscopy is not a static technique; it has continued to evolve. Let's look at some important extensions and future directions, starting with Fast-Ion-Beam Velocity Modulation.

This is a very sophisticated variant of the technique. As the first bullet point states, instead of a glow discharge cell, you start with a well-collimated beam of ions that has been created in a source, mass-selected, and then accelerated to high kinetic energies—typically on the order of 10 4 eV 10^4 eV—in a high-vacuum chamber. This gives you a pure beam of a single, known ion.

How do you apply the velocity modulation? As the second bullet explains, you superimpose a radio-frequency, or RF, electric field *perpendicular* to the direction of the ion beam's travel. This wiggles the ions' trajectory

slightly, but more importantly, it imposes a small velocity modulation along the laser axis without disturbing the overall beam alignment.

The result, as the final bullet highlights, is a technique that enables Doppler-free spectroscopy. Because all the ions in the beam have nearly the same velocity, the inhomogeneous Doppler broadening is dramatically reduced. This allows for incredible, sub-Megahertz resolution. We can essentially perform VMS on a mass-selected, velocity-selected sample. This is an ultra-high resolution technique, which is discussed in more detail in other sections of the course.

Page 41:

Here are a couple of other powerful extensions that enhance the sensitivity and specificity of VMS.

First, is the combination with Cavity Ring-Down Spectroscopy, or CRDS. We've talked about using a multipass cell to increase the path length. An optical cavity takes this to the extreme. By integrating the VMS discharge *inside* a high-finesse optical cavity, formed by two highly reflective mirrors, the effective path length can be amplified by factors of 10,000 or even more. This incredible enhancement in path length permits the detection of extremely weak absorption signals, such as the very faint overtone vibrational transitions of molecular cations. This technique is often called Cavity-Enhanced VMS, or CE-VMS.

Second, there are multi-frequency lock-in techniques. This involves modulating more than one thing at a time. For example, one could simultaneously modulate the discharge velocity at a frequency f 1 f_1, and

also modulate the laser intensity or frequency at a different frequency, f 2 f_2. By using two lock-in amplifiers or a dual-phase lock-in capable of two-dimensional demodulation, one can separate signals that arise from different physical effects. This is particularly useful for separating pure velocity modulation effects from non-linear saturation effects that can sometimes occur with high laser powers, leading to even cleaner and more reliable spectra.

Page 42:

This final slide provides an excellent flowchart summarizing the advanced VMS variants we've just discussed, highlighting their key features and performance.

On the left, we have **Fast-Ion-Beam Velocity Modulation**. The diagram shows an ion source, an extraction and acceleration stage, the perpendicular RF field for modulation, a laser beam intersecting the ion beam in a high-vacuum chamber, and a detector. The key achievable performance characteristics are sub-MHz, Doppler-free resolution, high sensitivity, and, critically, mass-selective ion selection. The key advantage is its ability to eliminate Doppler broadening and isolate specific ions for ultra-high-resolution studies.

In the center is **Cavity-Enhanced VMS (CE-VMS)**. The diagram shows the VMS discharge placed *inside* an optical cavity formed by two high-reflectivity mirrors. The achievable performance has standard VMS resolution, but with ultra-high sensitivity. The path length is amplified by factors up to 10 4 10^4 or more. The key advantage is its extreme

sensitivity, which enables the detection of very weak transitions like overtones.

On the right is **Dual-Modulation VMS**. The diagram shows two modulation sources: one at frequency f 1 f_1 for the discharge, and another at frequency f 2 f_2 for the laser intensity. The signals are processed using two-dimensional demodulation. The performance offers standard VMS resolution and sensitivity, but with very high signal purity. The key advantage here is its ability to isolate pure velocity-modulation signals from any artifacts that might arise from laser amplitude fluctuations or saturation effects.

These advanced techniques demonstrate the continued vitality and adaptability of the core VMS concept, pushing the frontiers of sensitivity and resolution in the spectroscopy of molecular ions. This concludes our lecture on Velocity-Modulation Spectroscopy. Thank you.