Vol. 2

Chapter *1.2 – 1.2.3*

Page 1:

Good morning, everyone. Welcome back to Phys 608, Laser Spectroscopy. I'm Distinguished Professor Dr M A Gondal, and Today, we begin a new and critically important chapter, Chapter 1.2, where we will explore the High-Sensitivity Methods of Absorption Spectroscopy.

In our previous discussions, we laid the theoretical groundwork for how light interacts with matter. But today, we bridge the gap between that fundamental theory and the cutting-edge reality of modern experimental physics. The techniques we will discuss in this chapter are the workhorses of research labs around the world, enabling scientists to detect minuscule traces of substances and observe phenomena that would otherwise be completely invisible.

This is where the true power and elegance of laser spectroscopy come to life. We're going to move beyond the simple picture of sending a laser through a sample and into the clever, sophisticated strategies that allow us to measure the seemingly unmeasurable. So, let’s begin.

Page 2:

Alright, let's start with the fundamental question: what is our motivation? Why do we need to seek "ultra-sensitive" methods for measuring absorption? Why isn't the basic experiment good enough?

The first bullet point gets to the heart of the matter. Many, if not most, real-world samples of interest produce absorption signals that are incredibly weak. Think about the challenges we face in science and technology.

An environmental scientist might need to detect a pollutant gas in the atmosphere at a concentration of parts-per-billion or even parts-per-trillion. An astrophysicist might be looking for the faint spectral signature of a molecule in a distant interstellar cloud. Or, a fundamental physicist might be searching for a so-called "forbidden" molecular transition—a transition that is incredibly unlikely to happen but whose observation could test the standard model of particle physics. In all these cases, the change in laser power due to the absorption is minuscule, often several orders of magnitude weaker than the inherent noise of even the best photodetectors. The signal is, quite literally, buried in the noise.

This leads directly to our second point, which frames the problem mathematically. Direct absorption measurements, which rely on the simple Beer–Lambert law, often fail when the product of the absorption coefficient, α α, and the path length, L L, is much, much less than 10 − 3 10^{-3}.

This quantity, α L αL, is the absorbance, or optical depth. A value of 10 − 3 10^{-3} means that only 0.1 % 0.1% of the light is absorbed as it passes through the sample. When the absorption is this weak, or weaker, trying to see that tiny dip in power is like trying to hear a single person whispering in the middle of a roaring stadium. The random fluctuations of the background noise will almost certainly drown it out. This is the central challenge that we must overcome.

Page 3:

So, let’s be more specific about what’s obscuring our signal. The tiny change in power we are trying to detect, the differential signal, is Δ P ΔP, which is the incident power, P 0 P_0, minus the transmitted power, P T P_T. This Δ P $\mathit{\Delta} P$ is our precious signal. The problem is that it gets swamped by two primary sources of noise: first, source power fluctuations—even the most stable laser has some inherent flicker or intensity noise, which we call δ P 0 δP_0. Second, there is the fundamental noise of the detector itself. These noise sources completely obscure the tiny Δ P $\mathit{\Delta} P$ we are looking for.

This defines the goal of all high-sensitivity techniques. Our objective is to design an experiment that can reliably detect fractional absorptions, which is the ratio of Δ P $\mathit{\Delta} P$ to P 0 P_0, down to fantastically small numbers. We're talking about detecting fractional absorptions of 10 − 7 10^{-7}—that's one part in ten million—all the way down to the truly mind-boggling regime of 10 − 17 10^{-17}. That's one part in ten quadrillion.

How on earth can we achieve this? We need a strategy. The first and most intuitive part of our strategy is stated here: we need to find a way to use very long *effective* path lengths, which we denote as L e f f L_eff. The logic is simple. The total absorption is proportional to the product α α times L L. If the absorption per unit length, α α, is tiny, we can still make the total absorption significant by making the path length, L L, enormous. If we can make our laser beam travel not just a few centimeters, but perhaps a few kilometers through the sample, our tiny signal will be greatly magnified. Of course, we can't build a kilometer-long laboratory, so we will need to find clever ways to fold that long path length into a very compact space.

Page 4

Building on the idea of increasing the path length, our overall strategy has two other pillars.

The second, and perhaps most powerful, strategy is to implement noise-rejection techniques. If we can't make our signal bigger, maybe we can make our noise smaller, or at least, make our measurement insensitive to it. The key idea here is *modulation*. Imagine again you're in a noisy room trying to hear a whisper. It's impossible. But if that person starts singing their message at a very high, specific pitch, your ear can easily pick it out from the low-frequency rumble of the crowd. This is exactly what we do in spectroscopy. We use techniques like frequency modulation or phase modulation to encode our absorption signal onto a high-frequency carrier wave. We then use a specialized electronic tool called a lock-in amplifier, which is tuned to that exact frequency, to demodulate the signal. This allows us to move our measurement away from the noisy, low-frequency world of laser flicker and into a pristine, high-frequency domain where the noise is fundamentally lower.

The third pillar of our strategy is careful optical mode engineering. This is really the practical implementation of the first two ideas. It's about how we design the physical hardware—the mirrors, the lenses, the cavities—to achieve our goals. This includes designing clever multipass cells, like the Herriott cells we'll see shortly, that use geometric optics to fold a long path length. It also includes using high-Quality-factor, or high-Q, optical cavities, which use resonance and wave optics to build up enormous light intensity and achieve long effective path lengths. And it includes advanced techniques like placing the sample directly inside the laser cavity itself to exploit the laser's own internal dynamics.

These three strategies—long path lengths, noise rejection via modulation, and clever optical engineering—form the complete toolkit for high-sensitivity spectroscopy.

Page 5:

This diagram provides a superb visual summary of the strategies we’ve just discussed and the dramatic improvements they provide. Let’s walk through it together.

On the far right, we have a vertical axis plotting the "Fractional Absorption" we can detect. It’s a logarithmic scale, ranging from 10 − 3 10^{-3} at the top, down to 10 − 13 10^{-13} at the bottom. Higher on this scale is worse sensitivity; lower is better.

Now, look at the top panel. This represents the simple, direct absorption experiment. We have a source, a sample cell, and a detector. The product α L αL is very small. The signal trace to its right shows the result: the tiny signal, Δ P ΔP, is completely lost in the thick, fuzzy band of noise. A direct measurement like this typically hits a wall at a sensitivity of about 10 − 3 10^{-3}, as indicated on the scale.

Next, let’s move to the middle panel. This illustrates our first strategy: increasing the path length. The diagram shows a multipass cell, where the laser beam, shown in red, bounces back and forth between two mirrors many times. This greatly increases the effective path length, Leff L_{eff}. Now look at the corresponding signal trace. Because the absorption has been magnified by the long path, the "Enhanced signal" is now clearly visible, poking out above the noise floor. This simple strategy can improve our sensitivity by several orders of magnitude, getting us down to the region of 10 − 7 10^{-7} on our scale.

Finally, the bottom panel demonstrates the power of noise rejection. Here, we've introduced a lock-in amplifier into our detection scheme. The key insight is shown in the graph on the right, which plots Power versus Frequency. Most technical noise, what we call "one over f f noise", is concentrated at low frequencies, or DC. This is shown by the pink shaded area, which is very high near zero frequency and falls off rapidly. A simple absorption measurement happens at DC, right in the heart of this noise. The genius of modulation, as shown in the diagram, is to move our signal away from DC to a high modulation frequency, f m f_m. Out at this high frequency, the technical noise is gone, and we are left with a much lower, flat noise floor, which we call the Shot Noise Limit. This is the fundamental quantum noise of light itself. By measuring in this quiet window, we can achieve sensitivities that are truly phenomenal, pushing down to 10 − 11 10^{-11}, 10 − 13 10^{-13}, or even lower, as shown on the scale.

Page 6:

Alright, now that we have the conceptual framework, let's put down the fundamental mathematical quantities we'll be working with, all of which stem from the Beer-Lambert law.

First, we have the incident light. We describe this by its spectral power density, denoted P 0 (ω) $P_0(\omega)$, or P-naught of omega. Let's be very careful about the units here: Watts per Hertz. This isn't just the total power of the laser; it's the power per unit frequency interval. This is essential for spectroscopy because our absorption features are, of course, frequency-dependent. So, P 0 (ω) $P_0(\omega)$ tells us how much power we have at the specific angular frequency ω ω.

Next, we have the transmitted power density. After the light travels a distance ' x x' through our absorbing medium, the power is attenuated. The transmitted power density, P T (ω) $P_\mathrm{T}(\omega)$, is given by the famous Beer-Lambert equation:

P T (ω) = P 0 (ω) e − α (ω) x .

$$P_\mathrm{T}(\omega) = P_0(\omega)\, e^{-\alpha(\omega)\, x}.$$

This equation describes the exponential decay of light as it passes through an absorber.

Now let's define the terms in that exponent, which is where the physics lies. This is highlighted under "Symbols and units".

The most important term is α (ω) $\alpha(\omega)$. This is the absorption coefficient. It is a fundamental property of the material itself and it quantifies how strongly the material absorbs light *at that specific frequency* ω ω. The fact that α α is a function of frequency is the entire basis of spectroscopy—the structure of α α versus ω ω is the unique spectral fingerprint of the atom or molecule we are studying. The standard unit for α α is inverse centimeters, or centimeters to the minus one (c m − 1 cm^{-1}). This can be interpreted as the fractional loss of intensity per centimeter of path length.

Page 7:

Continuing with our definitions, the symbol x x in the Beer-Lambert equation represents the physical path length that the light travels through the sample. Its unit is typically centimeters, to be consistent with the units of the absorption coefficient. The symbol ω ω, as we've said, is the angular frequency of the light, equal to 2 π v $2\pi\nu$, and its unit is radians per second.

Now we come to a critically important tool that we will use repeatedly throughout this course: the small-signal approximation. In all the high-sensitivity scenarios we are interested in, the total absorption is extremely weak. This means the exponent in the Beer-Lambert law, the product α (ω) x $\alpha(\omega)x$, is a very small, dimensionless number, much less than one. When this condition holds, we can use the first-order Taylor series expansion for the exponential function, which is that e − y ≈ 1 − y $e^{-y} \approx 1 - y$, for small y y.

Applying this to the Beer-Lambert law gives us a much simpler, linear relationship: The transmitted power, P T P_{T}, is approximately equal to the incident power, P 0 P_0, multiplied by the quantity, in brackets, [1 − α (ω) x] $[1 - \alpha(\omega)x]$.

This approximation makes it very easy to calculate the quantity we are actually trying to measure: the differential signal, capital Delta P. This is simply the light that was lost, which is the incident power P 0 P_0 minus the transmitted power P T P_{T}. Using our small-signal approximation, this becomes:

Δ P (ω) = P 0 − P T ≈ P 0 α (ω) x .

$$\mathit{\Delta} P(\omega) = P_0 - P_{\mathrm{T}} \approx P_0\,\alpha(\omega)x.$$

This simple, linear relationship is beautiful. It tells us that the signal we hope to measure is directly proportional to three things: the power of our light source, P 0 P_0; the strength of the absorption, α α; and the path length through the sample, x x. This makes our strategies immediately obvious: to get a bigger signal, we should use a brighter laser, a longer path length, or find a stronger transition.

Page 8:

Now, let's connect the macroscopic absorption coefficient, α α, to the microscopic world of atoms and molecules. This brings us to the minimum detectable number density.

For a dilute gas sample, which is often the case in high-sensitivity measurements, the absorption coefficient α α is not just an abstract number; it's directly proportional to how many absorbing molecules we have. The relationship is given by the first equation on this slide:

α = N 1 σ i k

$$\alpha = N_1 \, \sigma_{ik}$$

Let's break this down. N 1 N_1 is the number density of our absorbing species, specifically, the number of molecules per unit volume that are in the correct lower energy state, state 'i', from which the absorption can occur. Its units are typically molecules per cubic centimeter.

The term σ i k σ_{ik} is the absorption cross-section for the transition from the initial state 'i' to the final state 'k'. You can think of this cross-section as the effective "target area" that each molecule presents to the incoming photons. It has units of area, typically square centimeters. A large cross-section means a strong, or "allowed," transition, while a small cross-section means a weak, or "forbidden," transition.

Now we can formulate the ultimate figure of merit for an entire spectroscopic instrument: what is the minimum number density of a substance, N 1 , min $N_{1,\min}$, that we can possibly detect? The detectability criterion incorporates all our instrumentation parameters into one powerful equation:

N 1 , min ≥ N E P a P 0 σ i k L ′

$$N_{1,\min} \geq \frac{\mathrm{NEP}}{a\, P_0\, \sigma_{ik}\, L'}$$

Here, N 1 , min $N_{1,\min}$ is our goal – we want to make it as small as possible. To do that, we need to understand the parameters on the right-hand side, which we'll define on the next page. This equation beautifully encapsulates the interplay between the properties of our sample (σ σ), our laser (P 0 P_0), our interaction length (L ′ L'), and our detection system (N E P NEP).

Page 9:

Let's now define the parameters from that crucial detectability criterion equation.

First, we have N E P NEP, which stands for Noise-Equivalent Power. This is a fundamental figure of merit for any photodetector. Its units are W / H z $\mathrm{W}/\sqrt{\mathrm{Hz}}$. The N E P NEP is formally defined as the optical signal power that produces a signal-to-noise ratio of one in a one-Hertz detection bandwidth. In simpler terms, it's the faintest whisper of light that the detector can distinguish from its own internal noise. A smaller N E P NEP means a better, more sensitive detector.

Next is the lowercase a a. This is a dimensionless factor, typically with a value close to one. It's a catch-all parameter that accounts for things like the overlap between the laser lineshape and the absorption lineshape, as well as the effective detection bandwidth of our system. For our strategic thinking, we can treat it as a constant of order unity.

L L is the single-pass interaction length through the sample, in centimeters.

With these definitions in hand, we can now clearly see our "strategy space" emerge directly from the mathematics of the N min N_{min} equation. To make N min N_{min} smaller, meaning to improve our sensitivity and detect fewer molecules, we must do one or more of the following things:

- First, we can increase P 0 P_0, the incident laser power. This is the brute force approach: just use a brighter laser. This appears in the denominator, so a larger P 0 P_0 makes N min N_{min} smaller. - Second, we can increase L L, the interaction length. This can be our physical length, or, as we've discussed, an effective length, L e f f L_{eff}, achieved using multipass cells or cavities. Again, this is in the denominator. - Third, we can decrease the N E P NEP of our detection system. This is in the numerator, so a smaller N E P NEP leads to a smaller N min N_{min}. This can be achieved by choosing a better detector with lower intrinsic noise, or by using clever techniques like lock-in detection which dramatically reduce the effective noise bandwidth, thereby reducing the effective N E P NEP of the entire system. - And of course, there's a fourth strategy, related to the sigma term: whenever possible, we should choose to probe the strongest possible transition of our target molecule, the one with the largest absorption cross-section, σ i k σ_{ik}.

Page 10:

Let's now dig deeper into exactly why direct absorption measurements fail for tiny values of the absorbance, α L αL. This slide presents a simple but devastating calculation.

As we've established, the signal we measure is the change in power, Δ P ΔP, which is the incident power P 0 P_0 minus the transmitted power P T P_T. And for small absorptions, we use the approximation that the transmitted power P T P_T is roughly P 0 P_0 times the quantity 1 − α L $1 - \alpha L$.

Now, consider the uncertainty, or error, in our measurement. The dominant source of uncertainty is usually the fluctuation in the laser power itself, which we call δ P 0 δP_0. The relative error in our *signal* is the uncertainty in the signal, δ (Δ P) $\delta(\Delta P)$, divided by the signal itself, Δ P ΔP. What is this ratio? Well, the uncertainty in our measurement, δ (Δ P) $\delta(\Delta P)$, is going to be driven primarily by the laser fluctuation, δ P 0 δP_0. The signal itself, Δ P ΔP, is approximately P 0 α L $P_0 \alpha L$.

So, the relative error in our measurement is approximately δ P 0 P 0 − α L $\frac{\delta P_0}{P_0 - \alpha L}$. We can regroup these terms to get the expression on the slide: The fractional error, δ (Δ P) Δ P $\frac{\delta(\Delta P)}{\Delta P}$, is approximately equal to the fractional laser noise, δ P 0 P 0 $\frac{\delta P_0}{P_0}$, multiplied by a huge amplification factor of 1 α L $\frac{1}{\alpha L}$.

This is the killer. The inherent fractional noise of our laser source is amplified by the reciprocal of our fractional absorption. If α L αL is very small, say 10 − 4 10^{-4}, then 1 α L $\frac{1}{\alpha L}$ is ten thousand. This means our relative measurement error is ten thousand times larger than the stability of our laser! Let's see what this means with a concrete example.

Page 11:

Let's plug in some realistic numbers to see how catastrophic this effect is.

Suppose we are trying to measure a very weak absorption where the absorbance, α L αL, is equal to 10 − 4 10^{-4}. This means only zero-point-zero-one percent of the light is absorbed.

Next, let's assume we have an excellent, research-grade laser with a power stability, Δ P 0 / P 0 $\Delta P_0/P_0$, of 10 − 4 10^{-4}. This means the laser power only flickers by about zero-point-zero-one percent. This is very good, but realistic.

Now, let's calculate the relative error in our signal using the formula from the previous page. The relative error, δ (Δ P) Δ P $\frac{\delta(\Delta P)}{\Delta P}$, is approximately the product of our laser stability and the amplification factor. That's 10 − 4 × 1 10 − 4 $10^{-4} \times \frac{1}{10^{-4}}$. The result is 1 1.

A relative error of one means that the uncertainty in our measurement is the same size as the measurement itself! Our signal is completely and utterly masked by the noise. It is impossible to make a reliable measurement under these conditions.

This leads us to two profound conclusions that will dictate the design of all advanced spectroscopic techniques. We need approaches that do two things:

- First, they must reduce our dependence on measuring the absolute power of the laser. We need a measurement that is immune to the fluctuations Δ P 0 ΔP_0.

- Second, they must find a way to convert the tiny absorption signal into a different, larger, and more easily separated observable quantity. For example, instead of measuring a tiny change in amplitude, perhaps we can measure a change in the light's phase, or its frequency, or the decay rate of light in a cavity. This is the path forward.

Page 12:

So, having established the problem and the general principles for a solution, this slide provides our road-map. It outlines the major categories of high-sensitivity schemes that we will explore.

First, we have the most direct approach: the enhancement of the effective path length, L. This directly attacks the 'L' in the alpha-L product to make the total absorption larger. There are two main families of techniques here. The first are geometric multipass cells, such as White cells or Herriott cells, which use mirrors to fold the beam path many times. The second are resonant cavities, specifically external Fabry-Pérot enhancement cavities, which use the principle of constructive interference to achieve long effective path lengths and build up very high optical power.

Second, we have modulation methodologies. This approach tackles the noise problem head-on. By modulating some property of the light or the sample, we shift our signal to a high frequency, away from the dominant low-frequency noise. The key techniques here are Frequency Modulation, or FM, and Phase Modulation, or PM, which are always used in conjunction with lock-in detection. We'll also touch upon more advanced variants like two-tone FM, which is a clever way to reap the benefits of very high-frequency modulation while using more conventional electronics.

Third, we have a particularly powerful and elegant technique known as Intracavity Laser Absorption Spectroscopy, or ICLAS. The core idea here is to place the absorbing sample *inside* the laser cavity itself. This allows the sample to interact with the enormous internal field of the laser. Furthermore, it leverages the sensitive, non-linear dynamics of the laser, such as mode competition, to produce a massively amplified signal. We'll see that this technique has several variations, including single-mode, multimode, and time-resolved methods.

Page 13:

The fourth category on our road-map consists of hybrid and derivative techniques. These are often the most advanced methods, as they cleverly combine principles from the other three categories to achieve the ultimate in sensitivity. Classic examples include cavity-ring-down spectroscopy, or

C-R-D-S, which is a time-domain method that measures the decay rate of light in a high-Q cavity, thereby becoming immune to laser power fluctuations. Another example, often considered the pinnacle of sensitivity, is a technique with a famously long acronym: NICE-OHMS, which stands for Noise-Immune Cavity-Enhanced Optical Heterodyne Molecular Spectroscopy. As its name implies, it combines a high-Q enhancement cavity with frequency modulation techniques to achieve truly astonishing detection limits.

So, as we explore this zoo of different techniques and acronyms, it’s crucial to remember the final bullet point on this slide. It is the unifying principle of this entire chapter. Every single one of these approaches, no matter how complex it seems, ultimately solves the dilemma of weak absorption, where α L ≪ 1 $\alpha L \ll 1$, in one of two fundamental ways: either by dramatically increasing the effective path length, L e f f L_{eff}, or by dramatically reducing the effective noise bandwidth through modulation. The very best techniques, of course, do both.

Page 14:

This flowchart provides a perfect visual summary of our entire discussion so far, laying out the strategies for overcoming the α L ≪ 1 $\alpha L \ll 1$ limit. Let's trace it from the top down.

At the very top, in the red box, is "The Challenge: Detecting Tiny Absorption." The fundamental issue is stated clearly below it: our tiny signal, Δ P ΔP, which is approximately P 0 α L $P_0 \alpha L$, is masked by the source noise, δ P 0 δP_0.

From this central problem, four main strategic branches emerge, corresponding to the road-map we just outlined.

Branch 1 is Path Length Enhancement. Its physical principle is simple: use multiple reflections to increase the effective L L. The examples given are Multipass Cells, like White and Herriott cells, and resonant Fabry–Pérot Cavities.

Branch 2 is Modulation Methodologies. The physical principle here is more subtle: it's derivative detection and noise rejection. By modulating, we convert the absorption signal into its derivative and move it to a quiet frequency region. The examples are Frequency Modulation (FM), Phase Modulation (PM), and Two-Tone FM.

Branch 3 is Intracavity techniques, or ICLAS. The physical principle here is to exploit the internal dynamics of the laser itself, namely mode competition and the massive amplification of the intracavity field. Examples include single-mode, multimode, and time-resolved ICLAS.

Finally, Branch 4 covers Hybrid and Time-Domain methods. The physical principle for many of these, like Cavity Ring-Down, is to measure an exponential decay in time, which is inherently insensitive to amplitude fluctuations. Examples include Cavity Ring-Down Spectroscopy (CRDS) and the ultimate hybrid, NICE-OHMS.

This chart is our guide. We will now proceed down the first branch and explore the details of path length enhancement.

Page 15

Alright, let's begin with our first strategy: enhancing the path length using multipass geometry. This slide gives a qualitative picture of how these devices, such as a Herriott cell, work.

The basic setup consists of two concave, highly-reflective mirrors, each with a radius of curvature r r. These mirrors are placed facing each other, separated by a distance d d that is approximately equal to their radius of curvature. This specific near-confocal or re-entrant configuration is key to their stable operation.

The laser beam is then injected into this cell, typically through a small hole bored into the center of the first mirror, mirror 1.

Because of the curvature of the mirrors, the beam doesn't simply reflect back on itself. Instead, it hits the second mirror at a slight angle, which directs it to a new spot on the first mirror, and so on. The beam bounces back and forth between the two mirrors, tracing out a specific pattern of spots on each mirror's surface. After a well-defined number of reflections, which we'll call q q, the optical design ensures that the beam path lands on the input hole of mirror 1 once again, but this time at a different angle, which allows it to exit the cell and travel to a detector.

The result is a dramatic increase in the path length. The total effective path length, L e f f L_{eff}, is simply the number of passes, q q, multiplied by the distance between the mirrors, L L.

L e f f = q L .

$$L_{\text{eff}} = q\,L.$$

For example, a cell that is only 50 cm 50 cm long but provides 100 100 passes gives an effective path length of 50 m 50 m. This is a simple and robust way to significantly boost our absorption signal. The angle between adjacent spots on the mirror surface, θ θ, is determined by the precise mirror separation, and as we'll see, this is a critical parameter.

Page 16:

The geometry of the spot pattern is governed by the equation shown here:

cos (θ) = 1 − r d ′

$$\cos(\theta) = 1 - \frac{r}{d'}$$

Here, θ θ is the angle that separates the successive reflection spots as they rotate around the mirror's center. ' r r' is the mirror's radius of curvature, and ' d ′ d'' is an effective distance related to the mirror separation ' d d'. The full derivation comes from ray- transfer matrix analysis, which we won't go into here, but the important takeaway is that by making very small adjustments to the physical distance ' d d' between the mirrors, we can change the angle θ θ, and therefore change the number of passes ' q q' and the specific pattern of spots.

This leads to the crucial practical point mentioned in the second bullet: we design the cell such that there are no overlapping spots. Why is this so important? We must remember that our laser beam is a coherent electromagnetic wave. If the beam spots from different passes were to overlap on the mirror surface, they would interfere with each other. This interference would create unwanted intensity patterns, known as etalon fringes or interference fringes. As we scan our laser's wavelength to measure a spectrum, these fringes would cause large, rolling oscillations in our baseline signal, which would completely swamp the tiny absorption signal we are trying to detect. By ensuring the spots are spatially distinct, we can treat the cell using simple geometric optics, effectively creating one very long beam path and avoiding these destructive interference effects.

Page 17:

This slide gives us a wonderful, clear visualization of a Herriott multipass cell, illustrating the principles we've just discussed.

On the left, we see a side-on, cross-sectional view of the cell. We have the two concave mirrors, M1 and M2. M1 has a small hole for the laser to enter and exit. The mirrors have a radius of curvature r r and are separated by a distance d d. The red lines trace the path of the laser beam as it is injected, bounces back and forth dozens of times, and finally exits. You can clearly see how a very long optical path is folded into a compact physical volume.

The real elegance of the design is shown in the diagram on the right, which depicts the spot pattern on the surface of one of the mirrors. The red dots represent the points where the laser beam strikes the mirror. Instead of a random mess, the spots form a well-defined, stable ellipse or circle. The label indicates that for this particular configuration, we have q = 26 $q = 26$ reflections. The angle θ θ, which we saw in the equation on the previous page, is the angle between adjacent spots as viewed from the center of the mirror. This stable, re-entrant pattern is the hallmark of a well-designed Herriott cell.

Page 18

Of course, nothing in physics is perfectly efficient. Bouncing a laser beam off mirrors dozens or hundreds of times comes with a cost in the form of optical loss. This slide begins to build the "loss budget" for a multipass cell.

First, let's consider the loss from the mirrors themselves, even with no absorbing sample present. Even the best mirrors are not one hundred percent reflective. Let's define the power reflectivity of each mirror as capital R, which is a dimensionless number very close to one, for example, 0.999. A single round-trip for the light involves two reflections, one off each mirror. The attenuation factor for a single round-trip due to mirror losses, A-sub-mirror, is given by the equation:

A m i r r o r = exp (− 2 (1 − R)) .

$$A_{\text{mirror}} = \exp\big(-2(1 - R)\big).$$

For reflectivities R very close to 1, this is approximately one minus two times (one minus R). So, the fractional loss per round trip is just two times the loss per mirror. This is the intrinsic, unavoidable loss of the empty cell.

Next, we include our sample. Let's say the sample fills the cell and has an absorption coefficient alpha. A half-pass is the distance L from one mirror to the other. In a single round-trip, the light travels a distance of 2L through the sample. Therefore, the attenuation factor due to sample absorption per round trip, A-sub-sample, is given by:

Asample = exp (− 2 α L) .

$$A_{\text{sample}} = \exp(-2\alpha L).$$

These are the fundamental loss components for each round trip the light makes inside the cell.

Page 19:

Now, let's combine these loss factors to find the total transmitted intensity after the light has completed its full journey. If the light makes 'q' round-trips inside the cell, we simply multiply the attenuation factors for each round-trip 'q' times. This means we add the terms in the exponent 'q' times. This gives us the final equation for the transmitted intensity, I-sub-T: I-sub-T equals I-naught times the exponential of, in brackets, minus two q times the quantity one minus R, minus two q alpha L.

I T = I 0 exp (− 2 q (1 − R) − 2 q α L)

$$I_{\mathrm{T}} = I_0 \exp(-2\,q(1-R)\ -\ 2\,q\alpha L)$$

Let's dissect the exponent. The first term, minus two q times one minus R, represents the total cumulative loss from all the mirror reflections. The second term, minus two q alpha L, represents the total cumulative absorption from the sample. Notice that the effective path length here is L-eff equals two q L.

Leff = 2 q L

$$L_{\text{eff}} = 2\,qL$$

Now let's compute the numerical examples to get a feel for the numbers.

First, consider an empty cell, where alpha-L is zero.

α L = 0

$$\alpha L = 0$$

We use excellent mirrors with R equals 0.99, and we have q equals 100 passes.

R = 0.99 , q = 100

$$R = 0.99, \quad q = 100$$

The ratio of transmitted to incident intensity, I-T over I-naught, is e to the power of minus two times one hundred times (one minus zero point nine nine), which is e to the power of minus two.

I T I 0 = e − 2 × 100 × (1 − 0.99) = e − 2

$$\frac{I_{\mathrm{T}}}{I_0} = e^{-2\times100\times(1-0.99)} = e^{-2}$$

This is approximately zero point one four. This means that even with no sample, we lose 86% of our light just due to the imperfections in our mirrors over 100 passes.

Now, let's add a weakly absorbing sample. The slide uses a value of alpha-L equals 0.01. Let's re-read the formula carefully. The absorption term is 2*q*alpha*L. This would be 2* 100 *0.01 = 2. This would make the exponent e to the minus (2+2) = e to the minus 4. The slide shows an exponent of minus (2 + 0.02). This implies that the product 2*q*alpha*L equals 0.02. This means alpha*L must be 0.0001, or 10^-4. So let's assume this was the intended value. In this case, the transmitted intensity I-T over I-naught is e to the power of minus 2.02, which is approximately 0.133. The absorption caused a drop in transmission from 14% to about 13.3%.

2 q α L = 2 × 100 × 0.01 = 2

$$2\,q\alpha L = 2 \times 100 \times 0.01 = 2$$

e − (2 + 2) = e − 4

$$e^{-(2+2)} = e^{-4}$$

2 q α L = 0.02 ⟹α L = 10 − 4

$$2\,q\alpha L = 0.02 \implies \alpha L = 10^{-4}$$

I T I 0 = e − 2.02 ≈ 0.133

$$\frac{I_{\mathrm{T}}}{I_0} = e^{-2.02} \approx 0.133$$

The most important conclusion from this analysis is highlighted in the final bullet point: High mirror reflectivity, R, is absolutely critical. If R were, say, 0.90 instead of 0.99, the term (1-R) would be ten times larger, and we would lose nearly all of our light in just a few bounces. The performance of a multipass cell is fundamentally limited by the quality of the mirror coatings.

Page 20:

We will now transition from the geometric optics of multipass cells to a more powerful, wave-optics approach for path length enhancement: the external Fabry-Pérot enhancement cavity.

A Fabry-Pérot cavity, in its simplest form, is an optical resonator consisting of two highly parallel, highly reflective mirrors facing each other. It's essentially a laser cavity, but without a gain medium inside. We define the properties of the mirrors by their reflectivity. It's important to distinguish between the amplitude reflectivity, denoted by a lowercase 'r', which describes the reflection of the electric field, and the power reflectivity, capital R, which describes the reflection of the optical intensity or power. The two are related by R = r 2 $R = r^2$.

Now, for a cavity to work its magic, a crucial condition must be met, as stated in the second bullet point: the incident single-mode laser must be "mode-matched" to the cavity. A laser beam has a specific spatial intensity profile, which is typically a Gaussian shape called the TEM-zero-zero mode. The optical cavity also has a set of spatial modes that it naturally supports. Mode-matching means using a set of lenses to carefully shape our input laser beam so that its size and curvature perfectly match the fundamental TEM-zero-zero mode of the cavity. If the modes are mismatched, most of the light will simply reflect off the front surface of the first mirror, and we won't get the power enhancement effect we're looking for.

When the laser is both mode-matched and its frequency is on resonance with the cavity, a remarkable buildup of power occurs inside. Let's look at the expression for this intracavity power.

Page 21:

When the conditions of mode-matching and frequency resonance are met, the steady-state power inside the cavity, P i n t P_{int}, can become much, much larger than the incident power, P 0 P_0. The relationship is stunningly simple:

P i n t = P 0 1 − R

$$P_{\text{int}} = \frac{P_0}{1 - R}$$

Here, capital R R is the power reflectivity of the cavity mirrors, which we assume for simplicity are identical. Let's appreciate what this means. If we use mirrors with a reflectivity R R of 0.99, then 1 − R $1 - R$ is 0.01. The intracavity power is then P 0 P_0 divided by 0.01, which is 100 times the incident power! This is an enormous amplification. The factor 1 / (1 − R) $1/(1 - R)$ is called the power enhancement factor of the cavity.

The derivation of this formula comes from summing the amplitudes of all the partial waves that are transmitted and reflected inside the cavity. On resonance, all the waves that have made multiple round trips interfere constructively with the incoming wave, leading to this dramatic power buildup. It's a classic geometric series problem.

This power enhancement only happens at very specific frequencies, known as the cavity eigenfrequencies or longitudinal modes. These are the frequencies for which an integer number of half-wavelengths fit perfectly between the two mirrors. The formula for these resonant frequencies is:

v m = m c 2 d

$$\nu_m = \frac{m\,c}{2\,d}$$

Here, ' v m ν_m' is the frequency of the m m-th mode, ' m m' is a large integer called the mode index, ' c c' is the speed of light, and ' d d' is the physical separation between the mirrors.

This creates a major practical challenge. These resonance peaks are extremely sharp for a high- R R cavity. To get the power enhancement, we must keep our laser frequency, v L ν_{L}, perfectly tuned to the peak of one of these modes. This requires an active feedback system, or a "lock". We constantly monitor the cavity transmission and use a fast actuator, like a piezo-electric transducer on one of the mirrors, to make tiny adjustments to the cavity length ' d d' to keep it locked on resonance. The gold-standard technique for generating the error signal for this lock is called the Pound-Drever-Hall, or PDH, technique, which, as we'll see, is intimately related to the FM spectroscopy methods we will discuss later.

Page 22:

This diagram shows a complete, practical setup of an External Fabry-Pérot Enhancement Cavity, putting together all the pieces we've just discussed. Let's trace the path of the light and the control signals.

Starting on the far left, we have the "Incident Laser" with power P 0 P_0.

The beam first passes through a set of "Mode-matching lenses." As we discussed, these are essential for shaping the laser’s T E M 00 TEM_{00} Gaussian beam to perfectly match the fundamental mode of the optical cavity.

The shaped beam then arrives at the cavity itself. The cavity is formed by two mirrors: an input coupler and a high-reflector mirror, labeled M2, separated by a distance d d.

If the system is locked, the power builds up inside the cavity, creating a strong "Gaussian beam waist" with an internal power Pint P_{int} that is much, much greater than the incident power P0 P_0. This is where we would place our sample to be interrogated by this intense field.

A very small fraction of the intracavity light leaks through the second mirror, M2. This light is directed to a "Photodiode (PD)," which serves as the pick-off for our locking system.

The electrical signal from this photodiode is fed into the "PDH Lock Electronics" box. This is the brain of the feedback loop. It processes the signal and generates a "Piezo Control Signal," which is an error signal that tells us if we are on or off resonance.

This control signal is then applied to a "PZT," or piezoelectric transducer, that is physically attached to mirror M2. The PZT expands or contracts in response to the voltage, making nanometer-scale adjustments to the cavity length d d. This keeps the cavity’s resonance peak perfectly locked to the laser’s frequency, ensuring maximum power enhancement at all times. This entire system is a closed-loop feedback circuit.

Page 23:

Let's now plug in some numbers to get a concrete feel for the incredible power enhancement that a Fabry-Pérot cavity can provide. This slide gives a simple numerical example.

First, the given parameters. Let's assume our mirrors have a power reflectivity, capital R R, equal to 0.99 0.99. These are high-quality, but commercially available, standard laser mirrors.

Second, let's assume our incident laser power, P 0 P_0, is 100 mW $100\,\mathrm{mW}$. This is a very modest power level, easily achievable with a common laboratory diode laser.

Now, we want to calculate the intracavity power, P int P_{int}, using the formula we learned: P int = P 0 1 − R . $P_{\text{int}} = \frac{P_0}{1-R}$.

Plugging in the numbers, we get: P int = 100 mW 1 − 0.99 . $P_{\text{int}} = \frac{100\,\mathrm{mW}}{1-0.99}$.

This simplifies to 100 mW 0.01 . $\frac{100\,\mathrm{mW}}{0.01}$.

The result is 10,000 mW $10{,}000\,\mathrm{mW}$, which is equal to 10 W . $10\,\mathrm{W}$.

Let's pause and appreciate this result. We started with a modest, 100 mW $100\,\mathrm{mW}$ laser beam. By passing it into a simple resonant cavity made of two mirrors, we have generated an internal optical field of 10 W $10\,\mathrm{W}$. That's a power enhancement factor of 100 100. This is the incredible power of constructive interference and resonance. And with even better mirrors, say R = 0.999 $R = 0.999$, the enhancement factor would be 1000 1000, giving us 100 W $100\,\mathrm{W}$ of intracavity power from the same input laser. This is the key to many high-sensitivity techniques.

Page 24:

Now, let's see what this enormous intracavity power does for our ability to detect a weak absorption. We'll continue the numerical example.

Let's include a sample inside our cavity. We'll assume the sample has a very weak absorption coefficient, α α, of 10 − 5 c m − 1 $10^{-5}\,\mathrm{cm}^{-1}$. And we'll assume the sample cell itself is very short, with a length L L of only 0.2 c m $0.2\,\mathrm{cm}$.

First, let's calculate the single-pass absorption factor for this sample, which the slide labels α L αL. This is the product of the absorption coefficient and the length: 10 − 5 c m − 1 × 0.2 c m

$$10^{-5}\,\mathrm{cm}^{-1} \times 0.2\,\mathrm{cm}$$

gives a dimensionless absorption factor of 2 × 10 − 6 2×10^{-6}. This is a tiny number, representing an absorption of just two parts per million on a single pass.

Now, let's calculate the effective minimum detectable absorption coefficient, which we'll call α min $\alpha_{\min}$. We can use a simplified formula: α min $\alpha_{\min}$ is approximately the Noise-Equivalent Power of our detector, N E P NEP, divided by the product of the power interrogating the sample and the sample length.

α min ≈ N E P P i n t L .

$$\alpha_{\min} \approx \frac{\mathrm{NEP}}{P_{\mathrm{int}}\,L}.$$

In our case, the power is the huge intracavity power, Pint P_{int}.

Let's plug in some typical values. A good detector might have an NEP NEP of about 10 nW $10\,\text{nW}$. From the previous slide, our Pint P_{int} is 10 W $10\,\text{W}$, and our sample length L L is 0.2 cm $0.2\,\text{cm}$.

So, α min α_{min} is approximately 10 nW $10\,\text{nW}$ divided by the product of 10 W $10\,\text{W}$ and 0.2 cm $0.2\,\text{cm}$. This works out to 10 × 10 − 9 W 2 W cm ,

$$\frac{10 \times 10^{-9}\,\text{W}}{2\,\text{W}\,\text{cm}},$$

which gives an α min α_{min} of 5 × 10 − 9 cm − 1 $5 \times 10^{-9}\,\text{cm}^{-1}$.

This is a remarkable sensitivity. To put it in perspective, what would our sensitivity have been in a simple single-pass experiment using the same 100 mW $100\,\text{mW}$ laser? We would use P 0 P_0 (0.1 W $0.1\,\text{W}$) instead of Pint P_{int} (10 W $10\,\text{W}$) in our calculation. This would give a sensitivity a hundred times worse. The cavity provides a huge boost.

The slide notes an improvement of approximately one thousand times over a single pass with the same P 0 P_0. This is because the full analysis shows that the cavity enhances not only the power but also the effective path length. The total enhancement is related to a quantity called the cavity finesse, which for these parameters is about 314 314. This, combined with the power enhancement, readily leads to an overall sensitivity improvement of a factor of 1000 1000 or more.

Page 25:

We now pivot from our first strategy of path- length enhancement to our second major strategy: noise rejection through modulation. We will begin with the foundational technique of Frequency Modulation, or FM, Spectroscopy.

The basic idea of FM spectroscopy is conceptually straightforward. Instead of using a laser with a fixed, stable frequency, we intentionally modulate, or "wobble," the laser's frequency back and forth in a sinusoidal pattern. This modulation is centered on the frequency we wish to probe.

The instantaneous angular frequency of our laser, omega- sub- L of t, is described by the equation:

ω L (t) = ω 0 + a sin (Ω t) .

$$\omega_\mathrm{L}(t) = \omega_0 + a\sin(\Omega t).$$

Let's carefully define these terms. ω 0 ω_0 is the center frequency of our laser, which we can tune slowly to scan across a spectral feature.

Lowercase ' a a' is the modulation index or modulation depth. It determines how far in frequency the laser deviates from the center frequency. Its units are radians per second.

Capital Ω Ω is the modulation angular frequency. This is the rate at which we wobble the frequency back and forth. Crucially, we choose Ω Ω to be a high frequency, typically in the kilohertz to megahertz range, for reasons we will see shortly.

The core principle is this: as the laser's frequency sweeps back and forth across an absorption feature, the transmitted power will be modulated in response. If the laser's center frequency is on the side of an absorption line, one side of the frequency swing will be absorbed more than the other. This creates an imbalance, which results in a component of the transmitted power oscillating at the modulation frequency, capital Ω Ω. This oscillating signal is what we will detect.

Page 26:

So, we've generated a tiny optical signal that's oscillating at a high frequency, Ω Ω. How do we detect it? The key is an instrument called a lock-in amplifier. A lock-in amplifier is a sophisticated electronic device that acts as an extremely narrow bandpass filter combined with a phase-sensitive detector. We provide it with a reference signal at our modulation frequency, Ω Ω. The lock-in then looks at the signal from our photodetector and extracts *only* the component that is oscillating at that exact frequency and with a specific phase. It powerfully rejects noise at all other frequencies.

Now, what is the nature of this signal that the lock-in amplifier measures? If we use a small modulation depth a a—that is, a modulation depth much smaller than the linewidth of our absorption feature—it can be shown that the detected signal is proportional to the *first derivative* of the absorption coefficient with respect to frequency, d α d ω $\frac{d\alpha}{d\omega}$. We have converted our absorption measurement into a slope measurement.

This provides two enormous benefits:

First, as we've motivated, it makes our measurement largely immune to broadband intensity noise. The dominant " 1 / f $1/f$" noise of the laser is concentrated at low frequencies, near DC. By modulating at a high frequency, Ω Ω, we move our signal far away from this noise, into a quiet spectral region.

Second, it provides excellent background cancellation. Things like absorption or reflection from cell windows, or slow drifts in laser power, are all very slow, low-frequency effects. Since our lock-in amplifier is only sensitive to signals at the high frequency Ω Ω, these slow background effects are automatically rejected, leaving us with a clean, often near-zero, baseline.

Page 27:

This third benefit is really the ultimate consequence and the primary goal of the first two. By moving our signal to a high-frequency region to escape technical noise, and by using a lock-in amplifier to reject out-of-band noise, frequency modulation spectroscopy enables us to achieve shot-noise-limited detection.

Let's remind ourselves what this means. Shot noise is the fundamental quantum noise that arises from the discrete, particle-like nature of photons. Even a perfectly stable, classical light beam isn't truly constant; it's a stream of random photon arrivals. This randomness sets an absolute minimum noise floor for any optical measurement, a limit dictated by the laws of quantum mechanics.

The statement here is that by using modulation techniques at frequencies in the kilohertz to megahertz range—or even higher—we can effectively eliminate all other "technical" sources of noise (like laser flicker, mechanical vibrations, electronic interference) to such a degree that the only significant noise source remaining is this fundamental shot noise. Reaching the shot noise limit is the holy grail of sensitive optical measurements, as it means your experiment is as sensitive as the laws of physics will allow for a given optical power.

Page 28:

This diagram provides a clear illustration of the principle of operation of FM spectroscopy. It consists of two panels, plotting frequency and power as a function of time.

Let's first look at the top panel. The vertical axis is Frequency, and the horizontal axis is Time. The red curve shows the instantaneous frequency of our laser, ωL ω_{L}. You can see it oscillating sinusoidally around a center frequency, ω0 ω_0, with a modulation depth of ' a a'. In the background, the broad, blue-shaded feature represents the absorption profile of our sample, α(ω) $\alpha(\omega)$, which has a characteristic linewidth, γ γ. As the red line oscillates, it samples different parts of this absorption profile.

Now, look at the bottom panel. The vertical axis here is Transmitted Power, Ptrans P_{trans}, and the horizontal axis is again Time. This panel shows the consequence of the frequency modulation. As the laser's frequency sweeps across the absorption feature, the amount of light transmitted through the sample changes. The blue curve shows this time-dependent transmitted power. Notice that the power dips whenever the instantaneous frequency passes through the region of high absorption. Because the frequency is being modulated sinusoidally, the transmitted power is also modulated in a periodic, though not perfectly sinusoidal, way.

This modulated signal, described by the equation P t r a n s (t) ∝ e − α (ω L (t)) $P_{\mathrm{trans}}(t) \propto e^{-\alpha\left(\omega_{\mathrm{L}}(t)\right)}$, is what our photodetector sees. It is this signal that we feed into the lock-in amplifier, which will then extract the component oscillating at our chosen modulation frequency.

Page 29:

To understand precisely where the derivative lineshape comes from, let’s walk through a brief Taylor- series derivation of the modulated signal.

First, we begin with the fundamental expression for the transmitted power, P T P_{T}, at the instantaneous laser frequency, ω L ω_{L}. This is simply the Beer–Lambert law:

P T (ω L) = P 0 exp [− α (ω L) L] .

$$P_{\mathrm{T}}(\omega_{\mathrm{L}}) = P_0 \exp[-\alpha(\omega_{\mathrm{L}})\,L].$$

The second step is the core of the derivation. We assume that our modulation depth, a a, is small compared to the features of our spectrum. This allows us to perform a Taylor- series expansion of the transmitted power function, P T (ω L) $P_{\mathrm{T}}(\omega_{\mathrm{L}})$, around the center frequency, ω 0 ω_0. Remember that the frequency excursion from the center is a sin (Ω t) $a\sin(\Omega t)$. The Taylor expansion states that the function is approximately equal to its value at the center point, plus the first derivative at the center point times the small displacement, plus 1 2 $\frac{1}{2}$ the second derivative times the displacement squared, and so on.

This gives us the expression on the slide:

P T (ω L) ≈ P T (ω 0) + d P T d ω | ω 0 a sin (Ω t) + 1 2 d 2 P T d ω 2 | ω 0 a 2 sin 2 (Ω t) + higher-order terms .

$$P_{\mathrm{T}}(\omega_{\mathrm{L}}) \approx P_{\mathrm{T}}(\omega_0) + \left.\frac{dP_{\mathrm{T}}}{d\omega}\right|_{\omega_0} a\sin(\Omega t) + \frac{1}{2}\left.\frac{d^2P_{\mathrm{T}}}{d\omega^2}\right|_{\omega_0} a^2\sin^2(\Omega t)$$
$$+ \text{higher-order terms.}$$

This expansion breaks down our complex modulated signal into a series of simpler components.

Page 30:

Continuing our derivation, we now look at the terms in our Taylor expansion.

The third step is to focus on the signal that our lock-in amplifier will detect when it's tuned to the fundamental modulation frequency, Ω Ω. This signal comes from the first-order term in the expansion, as it is the only one that oscillates purely at sin (Ω t) $\sin(\Omega t)$. We can therefore write that the time-varying part of our transmitted power, Δ P T (t) $\Delta P_{\mathrm{T}}(t)$, is approximately equal to:

the derivative, d P T d ω $\frac{dP_{\mathrm{T}}}{d\omega}$, evaluated at ω 0 ω_0, multiplied by a sin (Ω t) $a\sin(\Omega t)$.

Δ P T (t) ≈ d P T d ω | ω 0 a sin (Ω t) .

$$\Delta P_{\mathrm{T}}(t) \approx \left.\frac{dP_{\mathrm{T}}}{d\omega}\right|_{\omega_0} a\sin(\Omega t).$$

This shows us that the amplitude of the signal at frequency Ω Ω is directly proportional to the slope, or derivative, of the transmitted power curve at our chosen center frequency.

The fourth step is to relate this derivative of the transmitted power back to the physical quantity we actually care about: the absorption coefficient, α α. We can do this using the small-signal approximation of the Beer-Lambert law, where P T ≈ P 0 (1 − α L) $P_\mathrm{T} \approx P_0(1 - \alpha L)$. We then take the derivative of this expression with respect to ω ω. P 0 P_0 and L L are constants, so we find that:

The derivative, d P T d ω $\frac{dP_\mathrm{T}}{d\omega}$, is approximately equal to − P 0 L d α d ω $-P_0 L \frac{d\alpha}{d\omega}$.

d P T d ω ≈ − P 0 L d α d ω .

$$\frac{dP_\mathrm{T}}{d\omega} \approx -P_0 L \frac{d\alpha}{d\omega}.$$

The slope of the power curve is directly proportional to the negative of the slope of the absorption curve.

Page 31:

Now we arrive at the final step, where we put everything together to find the output of our lock-in amplifier.

The lock-in amplifier measures the amplitude of the signal component that oscillates at the modulation frequency, Ω Ω. From step 3, we know this amplitude is 'a' times the derivative d P T d ω $\frac{dP_\mathrm{T}}{d\omega}$. From step 4, we know that d P T d ω $\frac{dP_T}{d\omega}$ is approximately − P 0 L d α d ω $-P_0 L\,\frac{d\alpha}{d\omega}$. However, the signal itself also depends on P 0 P_0. Let's revisit the previous slide. Ah, the derivation on the slide is slightly simplified. A more rigorous derivation would show that the final lock-in signal S (Ω) $S(\Omega)$ is proportional to P 0 a L d α d ω $P_0\,a\,L\,\frac{d\alpha}{d\omega}$. Let's assume the lowercase 'a' in the final equation on this slide is a combined constant. The key result is that the signal, S (Ω) $S(\Omega)$, is proportional to:

S (Ω) ∝ − a L d α d ω | ω = ω 0 .

$$S(\Omega) \propto -\,a\,L\,\frac{d\alpha}{d\omega}\bigg|_{\omega=\omega_0}.$$

This is the central result of small-modulation FM spectroscopy. It states that the signal measured by the lock-in amplifier is not the absorption profile itself, but rather its first derivative.

Thus, the crucial conclusion is that FM detection mathematically converts an absorption profile into a first derivative lineshape. In doing so, it has shifted our signal from DC to a high frequency, Ω Ω, giving us all the powerful noise-rejection benefits that we have discussed.

Page 32:

What about the higher-order terms in that Taylor expansion, like the one with a 2 sin 2 (Ω t) $a^2\sin^2(\Omega t)$? They don't just disappear; they produce signals at higher harmonics of the modulation frequency. This leads to a powerful set of techniques broadly known as "Derivative Spectroscopy."

A full Fourier analysis of the modulated signal reveals that it contains frequency components at all integer multiples of the fundamental modulation frequency, Ω Ω. That is, at n Ω $n\Omega$, where n n can be one, two, three, and so on.

We can use our lock-in amplifier to detect the signal at any of these harmonics, not just the fundamental. The general rules for what we measure are as follows, assuming a small modulation depth:

- If we detect at an *odd* harmonic, where n n is 1, 3, 5, and so on, the signal will be proportional to the nth *odd* derivative of the absorption profile, d n α d ω n $\frac{d^n\alpha}{d\omega^n}$. These signals will be carried on sine terms, meaning they are anti-symmetric about the line center. - If we detect at an *even* harmonic, where n n is 2, 4, 6, and so on, the signal will be proportional to the nth *even* derivative. These signals will be carried on cosine terms, making them symmetric about the line center.

This gives us an incredible tool: by simply changing the detection frequency of our lock-in amplifier, we can experimentally measure the first, second, third, or even higher derivatives of our spectral line, each of which reveals different information and has different advantages for analysis.

Page 33:

This slide presents the explicit mathematical forms for the signals at the first three harmonics, which are the most commonly used in practice. Let's examine them one by one.

First, the signal at the fundamental frequency, Ω $\mathit{\Omega}$, which we denote S (Ω) $S(\mathit{\Omega})$. This is the first harmonic, or 1-f signal. Its time-dependent form is:

S (Ω) = − a L d α d ω sin (Ω t)

$$S(\mathit{\Omega}) = -\,a\,L\,\frac{d\alpha}{d\omega}\,\sin(\mathit{\Omega} t)$$

The amplitude of this signal is proportional to the first derivative.

Next, the signal at the second harmonic, 2 Ω $2\mathit{\Omega}$, which we call the 2-f signal. Its form is:

S (2 Ω) = + a 2 L 4 d 2 α d ω 2 cos (2 Ω t)

$$S(2\mathit{\Omega}) = +\,\frac{a^2\,L}{4}\,\frac{d^2\alpha}{d\omega^2}\,\cos(2\mathit{\Omega} t)$$

The amplitude here is proportional to the second derivative of the absorption profile. Notice the prefactor contains an 'a-squared', which means for small modulation depths 'a', the second harmonic signal is generally smaller than the first.

Finally, the signal at the third harmonic, 3 Ω $3\mathit{\Omega}$, or the 3-f signal. Its form is:

S (3 Ω) = + a 3 L 24 d 3 α d ω 3 sin (3 Ω t)

$$S(3\mathit{\Omega}) = +\,\frac{a^3\,L}{24}\,\frac{d^3\alpha}{d\omega^3}\,\sin(3\mathit{\Omega} t)$$

This signal is proportional to the third derivative, and its amplitude is even smaller, scaling as 'a-cubed'.

These results show how we can directly access the different derivatives of our lineshape by tuning our detection electronics.

Page 34:

So, why would we want to measure these different harmonics? They have distinct shapes and properties that make them useful for different practical applications.

The first derivative signal, obtained by detecting at n = 1 $n = 1$, has that characteristic dispersive or 'S'-shape. Its most important feature is that it passes through zero exactly at the peak of the original absorption line. This zero-crossing point provides a very sharp, unambiguous marker for the line center. Furthermore, the signal around this zero-crossing is very steep, making it an ideal "error signal" for frequency stabilization. If the laser frequency drifts off the peak, a positive or negative voltage is generated, which can be fed back to the laser to correct its frequency. This is the fundamental principle behind many laser locking schemes, including the Pound-Drever-Hall technique.

The second derivative signal, obtained by detecting at n = 2 $n = 2$, has a very different character. It yields a symmetric peak that is centered precisely at the absorption maximum. Unlike the original absorption profile, however, this 2 f $2f$ signal has a natural zero baseline away from the resonance. The height of this central peak is directly proportional to the absorber concentration. This combination of a sharp central feature and a flat, zero baseline makes the second derivative signal excellent for quantitative analysis—for accurately measuring how much of a substance is present. It's one of the most widely used techniques in trace gas sensing.

Page 35:

These two graphs provide a perfect visual comparison between the standard absorption profile and the first harmonic signal.

In the top plot, we see the absorption coefficient, α (ω) $\alpha(\omega)$, plotted against frequency. It's a simple, symmetric peak, perhaps a Lorentzian or a Gaussian lineshape. The peak of the absorption is marked by the vertical dashed red line.

In the bottom plot, we see the corresponding first harmonic signal, S (Ω) $S(\Omega)$, which is proportional to the negative of the first derivative, − d α d ω $-\frac{d\alpha}{d\omega}$. As predicted by the mathematics, it has a dispersive, ‘S’-like shape. The most crucial feature to observe is that it passes exactly through zero at the precise frequency where the absorption is maximum. The steep, linear slope around this zero-crossing is what makes it such a superb discriminant for frequency locking applications.

Page 36:

This slide continues our visual exploration, showing the shapes of the second and third harmonic signals. The horizontal axis is labeled as "Frequency Detuning," which is the frequency (ω − ω 0) / Γ $(\omega - \omega_0)/\Gamma$. This allows us to see the universal shapes of these derivatives.

Let's look at the top plot, which shows the second harmonic signal, proportional to the second derivative of alpha. As we discussed, this lineshape is symmetric. It has a strong, negative-going central peak that occurs exactly at the line center, at zero detuning. This central peak is flanked by two smaller, positive-going lobes. The sharp central feature on a zero baseline makes this ideal for quantitative measurements.

Now, let's examine the bottom plot, showing the third harmonic signal, which is proportional to the third derivative. Like the first harmonic, this lineshape is anti-symmetric. It has a zero-crossing at the line center, but it exhibits more complex structure. It has a central positive lobe and two negative side lobes. While less commonly used than the first and second harmonics, the third derivative can sometimes be useful for resolving closely spaced or overlapping spectral features.

Page 37:

Now that we understand the power of frequency modulation, let's turn to the practical question of how we actually implement it. How do we make a laser's frequency wobble? This slide outlines the first of several common techniques.

Technique 1 is resonator length modulation. This method is often used for gas lasers, like Helium-Neon lasers, or for lasers that use an external cavity for tuning, such as external cavity diode lasers.

The implementation is conceptually simple: we mount one of the laser's end mirrors on a piezoelectric actuator, or PZT. A PZT is a ceramic material that physically expands or contracts when a voltage is applied across it. By applying a sinusoidal voltage to the PZT, we can make the mirror vibrate back and forth, thus modulating the physical length d d of the laser cavity.

Since the laser's output frequency is determined by the cavity length (recall that v = q c 2 d $\nu = \frac{qc}{2d}$), modulating the length d d directly modulates the frequency v ν. The relationship between the change in frequency, Δ v $\Delta\nu$, and the change in cavity length, Δ d Δd, is given by the equation:

Δ v = − q c 2 d 2 Δ d

$$\Delta\nu = -\frac{qc}{2d^2}\,\Delta d$$

This expression comes directly from differentiating the laser cavity resonance condition.

The major limitation of this technique is speed. A PZT is a mechanical object; it has mass and inertia. You cannot oscillate it at extremely high frequencies. It has its own mechanical resonances that limit its useful operating range. Typically, this method is restricted to modulation frequencies, Ω Ω, of at most a few kilohertz. This is often good enough to move away from the worst of the 1 / f $1/f$ noise, but it may not be fast enough to reach the true shot-noise limit.

Page 38:

For higher modulation frequencies, we need to employ faster, non-mechanical methods. This slide introduces two more powerful techniques.

Technique 2 is direct injection current modulation, which is primarily used for semiconductor diode lasers. The physics of a diode laser is such that its output frequency is dependent on the density of charge carriers in the semiconductor junction, which in turn is controlled by the injection current. Therefore, a very simple way to modulate the frequency is to add a small, high-frequency AC signal directly on top of the main DC drive current. This can be done with a simple function generator and a bias-T circuit. This method is incredibly effective and can achieve modulation frequencies up to hundreds of megahertz or even into the gigahertz range. The one significant drawback is that modulating the current also tends to modulate the laser's output *power* or amplitude. This unwanted amplitude modulation, known as residual amplitude modulation or RAM, can create spurious background signals that may limit sensitivity.

Technique 3 is the cleanest and most versatile method: using an external electro-optic phase modulator, or EOM. An EOM is a special crystal, such as lithium niobate, whose refractive index changes in response to an applied electric field (this is the Pockels effect). We take the stable, single-frequency output from our laser and pass it through this crystal. We then apply a high-frequency sinusoidal voltage across the crystal. This causes the crystal's refractive index 'n' to oscillate according to the equation:

n (V) = n 0 (1 + b sin (Ω t)) .

$$n(V) = n_0\big(1 + b\sin(\Omega t)\big).$$

This oscillation in the refractive index modulates the optical path length through the crystal, which imposes a direct sinusoidal modulation on the *phase* of the light wave. This generates a phase deviation, Δ ϕ m $\Delta\phi_{\mathrm{m}}$. As we will see, a pure phase modulation is extremely desirable and can operate at very high frequencies, well into the GHz range, with minimal unwanted amplitude modulation.

Page 39:

When we use an EOM to impose a phase modulation on our laser beam, we are doing something more profound than just "wobbling" the frequency. In the frequency domain, this action creates a whole family of new frequency components called sidebands. These sidebands appear at frequencies of ω 0 ± n Ω $\omega_0 \pm n\Omega$. So, we have sidebands at ω 0 ± Ω , ω 0 ± 2 Ω , … $\omega_0 \pm \Omega, \omega_0 \pm 2\Omega, \ldots$. The amplitudes of these sidebands are described by Bessel functions of the first kind, J q J_q, where the argument is the phase modulation depth, Δ ϕ m $\Delta\phi_\mathrm{m}$.

Now we arrive at a crucial distinction, as highlighted by the title of this slide: Phase vs. Amplitude Modulation, and its role in Noise Cancellation.

Let's first consider Amplitude Modulation, or AM. In pure AM, the modulation process creates a carrier and just a single pair of sidebands, at ω 0 ± Ω $\omega_0 \pm \Omega$. Crucially, in AM, these sidebands are created *in-phase* with the carrier wave.

Now consider Phase Modulation, or PM, which is what an ideal EOM produces. PM also creates sidebands. For small modulation depths, the dominant sidebands are also at ω 0 ± Ω $\omega_0 \pm \Omega$. However, there is a critical difference: PM yields symmetric sidebands that have *opposite phase*. That is, the electric field of the upper sideband (at ω 0 + Ω $\omega_0 + \Omega$) is 180 ∘ 180° out of phase with the electric field of the lower sideband (at ω 0 − Ω $\omega_0 - \Omega$). This specific phase relationship is the key to the powerful noise cancellation properties of FM spectroscopy.

Page 40: How does this opposite phase relationship in phase modulation lead to such effective noise cancellation? Let's walk through the process.

Our photodetector is a square-law detector; it measures intensity, which is the square of the electric field. When our phase-modulated light hits the detector, several different frequency components beat against each other. The lock-in amplifier, tuned to the modulation frequency Ω Ω, is sensitive to the beat note created between the strong central carrier and the two sidebands.

Specifically, the carrier beats with the upper sideband, producing an electrical signal at their difference frequency, which is Ω Ω. At the same time, the carrier beats with the lower sideband, also producing an electrical signal at their difference frequency, Ω Ω.

Now, consider the case without any absorption. The two sidebands have equal amplitude. Because their optical phases are opposite, the two electrical beat notes they produce at frequency Ω Ω are also 180 degrees out of phase with each other. Since they have equal amplitude and opposite phase, they sum together and perfectly cancel each other out. The result is a zero baseline signal.

Now, let's introduce our sample, which absorbs some of the light at, say, the upper sideband's frequency. This reduces the amplitude of the upper sideband. Now, the two beat notes are no longer equal in magnitude. They no longer cancel perfectly. This imbalance creates a net, non-zero signal at the lock-in frequency Ω Ω. This imbalance *is* our absorption signal.

Here is the magic of the noise cancellation: any fluctuation in the intensity of the main laser is a "common-mode" fluctuation. It affects the power of the carrier and both sidebands equally and in phase. Since our detection method is sensitive only to the *imbalance* or *difference* between the sidebands, this common-mode noise is automatically subtracted out and cancels to first order.

This is what provides the dramatic improvement in sensitivity.

Page 41:

This diagram provides an excellent visual explanation of the concepts we just discussed, comparing Amplitude Modulation on the left with Phase Modulation on the right.

Let's start with the top row, which shows the case with no absorption.

In panel (a), for AM, we see the frequency spectrum consists of the strong carrier at ω 0 ω_0 and two smaller sidebands at ω 0 ± Ω $\omega_0 \pm \Omega$. Crucially, they are all drawn as upward arrows, indicating they are in-phase.

In panel (b), for PM, we see the carrier and two sidebands. But here, the lower sideband is drawn as a downward arrow, representing its opposite phase. Below this, the process at the detector is illustrated. The beat note between the carrier and the upper sideband, Beat(C,USB) $\mathrm{Beat}(C,\mathrm{USB})$, is a sine wave. The beat note with the lower sideband, Beat(C,LSB) $\mathrm{Beat}(C,\mathrm{LSB})$, is a sine wave that is perfectly out of phase. When these two are summed (indicated by the Σ Σ symbol), the result is a flat line: "Perfect Cancellation → → Zero Baseline Signal."

Now, let's look at the bottom row, which shows what happens when there is absorption of the upper sideband.

In panel (a) for AM, the upper sideband's amplitude is reduced. This imbalance in the total power is detected directly as an AM signal, which is highly susceptible to laser intensity noise.

In panel (b) for PM, this is the key case. The absorption reduces the amplitude of the upper sideband. Now, the two beat notes are no longer equal in magnitude. When they are summed, they no longer cancel. A net sinusoidal signal at frequency Ω Ω appears. The diagram shows: "Imbalance → → Non-Zero Signal at Ω Ω". This signal, born out of an imbalance on a zero background, is our highly sensitive, noise-immune measurement.

Page 42:

We've established that modulating at a high frequency, capital Ω Ω, is crucial for noise rejection. But how do we choose the optimal value for Ω Ω?

First, we must remember that technical noise sources, like laser flicker and 1 / f $1/f$ electronic noise, are most severe at low frequencies and decrease rapidly as frequency increases. Therefore, the higher we can make Ω Ω, the better our noise rejection will be.

The optimum region for Ω Ω also depends on the properties of the spectral line we are trying to measure. A good rule of thumb is given by the inequality on the slide:

Γ ≤ Ω ≤ 10 Γ

$$\Gamma \leq \Omega \leq 10\,\Gamma$$

The modulation frequency Ω Ω, in units of radians per second, should be greater than or equal to the linewidth, Γ Γ, and less than or equal to about ten times the linewidth. In terms of ordinary frequency in Hertz, this means the modulation frequency should be comparable to or a few times larger than the linewidth of the transition (Γ / (2 π) $\Gamma/(2\pi)$). Modulating much slower than the linewidth means you're not really measuring a derivative, while modulating much faster can sometimes reduce the signal magnitude.

However, if ultimate sensitivity is the goal, we can push Ω Ω to be very high, greater than one gigahertz. At these very high frequencies, we are far beyond the reach of almost all technical noise sources. This allows us to suppress all technical noise and reach the fundamental shot-noise limit.

Page 43:

Operating at very high modulation frequencies in the gigahertz range presents a practical challenge: most standard laboratory photodetectors and lock-in amplifiers are not designed to work at such high frequencies.

To get around this, a clever technique called "two-tone FM" spectroscopy is often employed. This method allows us to get the noise-suppression benefits of GHz-frequency sidebands while performing the final electronic detection at a more manageable frequency, typically in the MHz range. It requires a down-conversion step before the lock-in amplifier, which we will touch on next.

So, with all these considerations, what is a good practical rule of thumb for choosing the modulation frequency Ω Ω? A very common and effective choice is to set the modulation frequency, capital Ω Ω, to be approximately equal to the full-width-at-half-maximum, or FWHM, of the target absorption line. This usually provides a good balance, yielding a strong derivative signal while ensuring the modulation is fast enough to provide significant noise rejection.

Page 44:

Let's briefly outline the principle of Two-Tone Frequency-Modulation Spectroscopy. This is an elegant solution to the problem of wanting to use GHz modulation frequencies without needing GHz detection electronics.

The process starts similarly to standard FM: a high-frequency carrier modulation is imposed on the laser beam using an EOM. Let's say we modulate at a few gigahertz. This creates the widely separated sidebands that are key to avoiding low-frequency laser noise.

Here's the trick: the radio-frequency signal that drives the EOM is not a pure, single-frequency sine wave. Instead, the gigahertz drive signal is itself amplitude-modulated at a much lower frequency, for example, in the megahertz range. This generates what are effectively two closely spaced FM carriers, or "tones." The result on the optical beam is a more complex pattern of sidebands, with two main clusters separated by the GHz frequency, and within each cluster, sidebands separated by the MHz frequency.

When this complex beam strikes the photodetector, a variety of beat notes are produced. Crucially, a beat note is generated between the two main tones. This beat note appears at the MHz difference frequency.

This MHz signal, which still carries the absorption information from the widely spaced GHz sidebands, is now at a frequency that is perfectly manageable for a standard lock-in amplifier.

The system still fully benefits from the gigahertz-scale separation of the sidebands from the carrier, which places them in the shot-noise-limited region, far from the 1/f noise of the laser.

Page 45:

The ultimate benefit of the two-tone FM technique is summarized here. It allows us to achieve the holy grail of shot-noise-limited detection, but it cleverly sidesteps the practical difficulty and expense of using gigahertz-bandwidth photodetectors and lock-in amplifiers in the demodulation chain.

It's a "best of both worlds" scenario. We get the supreme noise immunity that comes from probing our sample with frequencies separated by gigahertz, while performing all of our sensitive electronic detection and signal processing in the much more convenient and accessible megahertz frequency range. This makes it a powerful and widely used technique for pushing the absolute limits of absorption spectroscopy.

Page 46:

We now move to our third major strategy for high-sensitivity measurements: Intracavity Laser Absorption Spectroscopy, or ICLAS. The core principle of ICLAS is as simple as it is powerful.

Instead of taking the output of a laser and passing it through an external sample cell, we place the sample directly *inside* the laser resonator itself, between the two laser mirrors.

To understand why this is so powerful, we must recall how a laser works. The light inside a laser cavity bounces back and forth, being amplified by the gain medium on each pass. The intracavity field intensity builds up until the round-trip gain provided by the medium exactly equals the total round-trip losses. These losses include the light transmitted through the output mirror and any absorption or scattering from components inside the cavity.

The crucial consequence is that each photon inside the cavity circulates many, many times before it eventually leaks out. A typical photon might make hundreds or even thousands of round trips. This means that if our sample is inside the cavity, it is interrogated by the light not once, but hundreds or thousands of times. This results in a huge effective path length, which dramatically enhances the absorption signal.

There are several ways to detect the effect of this intracavity absorption. The slide lists four key detection concepts. The first is to simply measure the fluorescence emitted from the sample. As the sample absorbs the very intense intracavity light, it will re-emit some of that energy as fluorescence, which we can detect from the side of the cell with a separate detector.

Page 47:

Let's continue with the other detection concepts for ICLAS.

The second method is to monitor the change in the laser’s output power, which we would measure with a detector, detector 2, placed outside the output coupler. If the intracavity sample introduces an absorption loss at the laser’s frequency, the overall loss of the cavity increases. The laser will then stabilize at a new, lower output power. As we will see, this effect can be highly non-linear and extremely sensitive, especially when the laser is operated near its lasing threshold.

The third method is to exploit the phenomenon of mode competition in multimode lasers. Many lasers, such as dye lasers or Ti:sapphire lasers, have a very broad gain profile and can oscillate on thousands of different longitudinal modes, or frequencies, simultaneously. If our sample introduces a narrow absorption loss that overlaps with just one of these modes, that mode will be at a disadvantage in the competition for gain against all the other modes. Its intensity can be drastically reduced, or even completely extinguished.

The fourth concept is a powerful time-resolved method. Instead of letting the laser reach a steady state, we use what’s called a "step-function pump." We turn the laser on very quickly and observe how the different modes build up from noise over a very short time. Modes that experience absorption will grow much more slowly than their neighbors. In this dynamic method, the effective path length, Leff L_{eff}, is simply the speed of light, c c, multiplied by the generation time, t t. This can lead to astonishingly long effective path lengths.

Page 48:

This schematic provides a clear block diagram of a generic Intracavity Laser Absorption Spectroscopy, or ICLAS, experiment.

At its core, we have a laser resonator, defined by a High Reflector (HR) mirror on the left and an Output Coupler (OC) mirror on the right.

Inside this resonator, we have two key components. The first is the Gain Medium, which provides the optical amplification that makes the laser lase. This gain medium is energized by an external Pump Source. The diagram notes that this pump source could be a step-function for time-resolved experiments.

The second, and most crucial, component inside the cavity is our Sample Cell. This contains the atoms or molecules we wish to study.

The Intracavity Laser Field, represented by the red beam, circulates back and forth, passing through both the gain medium and the sample cell on every round trip. As we've discussed, this leads to an enormous Effective Path Length. In the time-resolved case, this is

Leff = ct.

$$L_{\text{eff}} = c\,t.$$

The diagram shows two possible ways to get a signal. Detector 1 is placed to the side of the sample cell and measures the fluorescence that is emitted as the sample absorbs the intense intracavity light. The second detection path is to measure the light that transmits through the Output Coupler. This measures the final output power of the laser, which, as we've discussed, is also sensitive to the intracavity absorption.

Page 49:

Let's now develop a more quantitative understanding of ICLAS, starting with the simplest case: a single-mode laser operating in a steady state. We'll perform a simple power balance analysis.

First, let's define the properties of our laser mirrors. The back mirror is a high reflector, so its reflectivity, R 1 R_1, is very close to one. The front mirror is the output coupler, which is designed to transmit a small fraction of the light. We'll call its transmission T 2 T_2. Therefore, its reflectivity, R 2 R_2, is 1 − T 2 $1 - T_2$.

Now, let's consider the relationship between the power inside the cavity, P int P_{int}, and the power that we measure coming out of the laser, P out P_{out}. The output power is simply the fraction T 2 T_2 of the intracavity power that hits the output mirror. Therefore, we can write:

P int = P out T 2 .

$$P_{\text{int}} = \frac{P_{\text{out}}}{T_2}.$$

We can define an intracavity power enhancement factor, q q, as being equal to 1 / T 2 $1/T_2$. This gives the simple relation:

P i n t = q P o u t .

$$P_{\text{int}} = q\, P_{\text{out}}.$$

For a typical laser with a 2% output coupler (T 2 = 0.02 $T_2 = 0.02$), this enhancement factor q q is 50. The power inside is 50 times the power outside.

Now, what is the total power absorbed by our sample per second? The absorbed power, Δ P (ω) $\Delta P(\omega)$, is the single-pass absorption factor, which is α (ω) $\alpha(\omega)$ times the sample length L L, multiplied by the power that is actually interrogating the sample. That power is the huge intracavity power, P i n t P_{int}.

Substituting our expression for P i n t P_{int}, we find:

Δ P (ω) = q α (ω) L P o u t .

$$\Delta P(\omega) = q\, \alpha(\omega)\, L\, P_{\text{out}}.$$

The power absorbed by the sample is enhanced by this same factor q q.

Page 50:

So, the power absorbed by the sample is enhanced by this factor q q, which is equal to one over the output coupler transmission, T 2 T_2. For typical lasers, T 2 T_2 can range from a few percent down to a fraction of a percent. This means the enhancement factor q q can readily be in the range of 50 to 500. This is a very significant enhancement.

This leads to several possible measurement modalities. We can aim to detect the absorbed power, Δ P ΔP, directly. Since this absorbed power is enhanced by the large factor q q, any detection method sensitive to absorbed energy will also have its sensitivity enhanced.

For example, the absorbed energy heats the sample. In a gas, this heating will cause a pressure rise. We can detect this pressure rise with a sensitive microphone. This is the basis of photo-acoustic spectroscopy, and placing the photo-acoustic cell inside a laser cavity is a well-known method for boosting its sensitivity.

Alternatively, as mentioned before, we can detect the fluorescence that is emitted after the molecule absorbs a photon. The amount of fluorescence will be proportional to the absorbed power, Δ P ΔP, and will therefore also be enhanced by the factor q q.

Page 51:

Let's now analyze the second detection method: measuring the change in the laser's output power. This approach is particularly sensitive when the laser is operated very close to its lasing threshold.

First, let's review the fundamental condition for steady-state laser operation. A laser reaches a stable output power when the gain it receives per round trip exactly balances the total losses it experiences per round trip. We call the saturated gain g-sub-s, and the total round-trip loss gamma. So, the threshold condition is:

g s = γ

$$g_s = \gamma$$

The gain of a laser medium is not constant; it saturates. For a homogeneously broadened gain medium, the saturated gain g-sub-s is given by the formula:

g s = g 0 1 + P P s

$$g_s = \frac{g_0}{1 + \frac{P}{P_s}}$$

Here, g 0 g_0 is the small-signal gain, which is the maximum gain available, determined by how hard we are pumping the laser. P P is the intracavity power, and P s P_s is the saturation power, which is a fundamental constant of the gain medium. This equation shows that as the power P P inside the laser builds up, the gain available for further amplification decreases.

The laser is a self-regulating system. The power P P will automatically adjust itself until the saturated gain g s g_s drops to a value that exactly equals the total loss, γ γ.

Page 52:

By setting the saturated gain equal to the loss (g s = γ $g_s = \gamma$) and solving the equation from the previous slide for the power P P, we can find the steady-state intracavity power of the laser:

P = P s (g 0 γ − 1) .

$$P = P_s\left(\frac{g_0}{\gamma} - 1\right).$$

Now, let's introduce our weakly absorbing sample into the cavity. This introduces a small additional loss, which we'll call Δ γ $\mathit{\Delta}\gamma$. This additional loss is equal to the absorption coefficient α α times the round-trip path length through the sample, which is 2 L $2\,L$:

Δ γ = α · 2 L .

$$\mathit{\Delta}\gamma = \alpha \cdot 2\,L.$$

This new loss will cause the laser power to drop to a new, lower value, P α P_α. We want to find the relative change in power, Δ P / P $\mathit{\Delta}P/P$.

The full derivation is a bit tedious, but it involves taking the derivative of the power expression with respect to the loss γ γ. The result, in the limit that Δ γ $\mathit{\Delta}\gamma$ is much smaller than γ γ, is shown on the slide:

Δ P P ≈ (g 0 g 0 − γ) 2 Δ γ γ .

$$\frac{\mathit{\Delta}P}{P} \approx \left(\frac{g_0}{g_0 - \gamma}\right)^2 \frac{\mathit{\Delta}\gamma}{\gamma}.$$

Actually, the formula on the slide is slightly simplified, the final expression for the sensitivity enhancement factor Q is the key part. Let's focus on that.

The response of the laser is highly non-linear. The sensitivity enhancement factor, Q Q, quantifies this non-linearity. It is the factor that multiplies our fractional added loss, Δ γ / γ $\mathit{\Delta}\gamma/\gamma$, to give our measured fractional power change, Δ P / P $\mathit{\Delta}P/P$. This enhancement factor is given by:

Q = g 0 γ (g 0 − γ) .

$$Q = \frac{g_0}{\gamma\,(g_0 - \gamma)}.$$

This expression holds the key to the high sensitivity of this method.

Page 53:

Let's take a close look at the behavior of this sensitivity enhancement factor, Q Q, in two different operating regimes.

First, let's consider the case where we are operating the laser far above its threshold. This means we are pumping it very hard, so the small-signal gain, g 0 g_0, is much, much larger than the round-trip loss, γ γ. In this limit, the term in the denominator, g 0 − γ $g_0 - \gamma$, is approximately equal to g 0 g_0. So, the expression for Q Q simplifies to:

Q ≈ g 0 γ g 0 .

$$Q \approx \frac{g_0}{\gamma\, g_0}.$$

The g 0 g_0s cancel, and we are left with Q Q is approximately 1 / γ $1/\gamma$. Since the loss γ γ is dominated by the output coupling T 2 T_2, this is essentially the same power enhancement factor ' q q' that we saw earlier. So, far above threshold, the enhancement is significant, but not extraordinary.

Now for the exciting case: when we operate the laser very close to its threshold. In this regime, we reduce the pump power so that the small-signal gain g 0 g_0 becomes just barely larger than the loss γ γ. This means the term in the denominator, g 0 − γ $g_0 - \gamma$, becomes a very, very small number. As this denominator term approaches zero, the enhancement factor Q Q becomes huge, approaching infinity in principle.

This suggests we could achieve unlimited sensitivity by operating precisely at the threshold. However, there is a catch. As we approach threshold, the laser becomes extremely sensitive not just to our sample's absorption, but to any small perturbation. It becomes highly susceptible to pump power noise, mechanical vibrations, and temperature fluctuations, making the output power very noisy and unstable. Therefore, in practice, there is always a trade-off between achieving maximum sensitivity and maintaining sufficient stability for a reliable measurement.

Page 54:

The analysis we just performed was for a laser with a homogeneously broadened gain medium, such as a dye, fiber, or solid-state laser. What happens in the case of a laser with an inhomogeneously broadened gain medium? A classic example is a gas laser, where the random thermal motion of the atoms leads to Doppler broadening.

In an inhomogeneously broadened medium, the gain saturates differently with power. The relationship for the saturated gain, g s g_{s}, is:

g s = g 0 1 + I I s

$$g_{\mathrm{s}} = \frac{g_0}{\sqrt{1 + \frac{I}{I_{\mathrm{s}}}}}$$

where I I is the intracavity intensity.

We can perform an analogous derivation to find the sensitivity to a small added loss. Without going through the detailed steps, the result for the fractional change in power is given on the slide:

Δ P P ≈ g 0 2 γ 2 Δ γ g 0 − γ ′

$$\frac{\Delta P}{P} \approx \frac{g_0^2}{\gamma^2} \frac{\Delta\gamma}{g_0 - \gamma'}$$

(Here γ ′ γ' is nearly identical to γ γ.)

The key feature to notice in this result is the even stronger dependence on operating near threshold. The enhancement factor now contains a term that looks like g 0 2 g_0^2, and the denominator still contains the g 0 − γ $g_0 - \gamma$ term that becomes very small near threshold. This even stronger non-linearity demonstrates that ICLAS can be even more sensitive when performed with an inhomogeneously broadened laser operating close to its threshold.

Page 55:

Now let's explore one of the most powerful ICLAS techniques: Multimode ICLAS, which relies on the phenomenon of mode competition.

First, let's set the scene. We are now considering a laser with a very broad gain profile, one that is wide enough to support lasing on a large number, N N, of longitudinal modes simultaneously. Imagine the laser output as a comb of thousands of discrete, sharp frequencies, all lasing at the same time under the same broad gain curve.

Now, what happens if we place a sample inside this laser's cavity, and that sample has a very narrow absorption line that happens to overlap with just one of these thousands of modes? Let's call it mode k k. This absorption introduces a small, additional loss, Δ γ $\Delta\gamma$, that is experienced *only* by mode k k.

The immediate effect is that the intensity of mode k k, I k I_k, will begin to drop. But that's not the whole story.

Because mode k k's intensity has dropped, it is now extracting less energy from the shared gain medium. This means the overall gain saturation is slightly reduced, and the gain medium "recovers" a little. This extra available gain is now up for grabs. It doesn't help mode k k, which is being actively suppressed by the absorption. Instead, this recovered gain is distributed among all the other, non-absorbing modes, boosting their intensities, I j I_j.

This process, where all the modes fight for a common pool of gain, is called mode competition. The degree to which they are coupled is quantified by a coupling parameter, K K, which ranges from 0 0 for no coupling to 1 1 for strong coupling. In many lasers, like dye lasers, this coupling is very strong, with K K close to one.

Page 56:

The consequence of this strong mode competition is a truly dramatic amplification of the effect of the absorption.

The relative power change for the single absorbing mode, mode k k, is given by the equation on this slide. Let's look at it closely. The first part of the expression, g 0 Δ γ γ (g 0 − γ) $\frac{g_0\Delta\gamma}{\gamma(g_0-\gamma)}$, is the same sensitivity enhancement factor we saw for the single-mode laser operating near threshold. But now, this is multiplied by a massive new factor: the quantity 1 + K N $1+KN$.

Let's think about this new factor. K K is the coupling strength, which is close to one. And N N is the number of competing modes, which can be in the thousands. This means the initial sensitivity enhancement is multiplied by an additional factor of several thousand.

In a laser with strong coupling (K ≈ 1 $K \approx 1$) and many modes, the result is an enormous contrast. Even a minuscule absorption on mode k k can cause a catastrophic loss of power for that mode. The absorbing mode can effectively vanish, as all of its energy is redistributed to the other modes. This creates what is known as a "spectral hole" in the laser's output spectrum.

The detection method is then elegant and straightforward. We take the broadband output from our multimode ICLAS laser and disperse it using a high-resolution monochromator or spectrograph with a detector array. The recorded spectrum will show the broad profile of the laser, but with sharp, dark lines appearing exactly at the frequencies where our intracavity sample absorbs. By observing the positions and depths of these missing mode intensities, we can map out the complete absorption spectrum with incredible sensitivity, often in a single shot.

Page 57:

Let's now turn to the final ICLAS variant we'll discuss: the powerful Time-Resolved ICLAS method, also known as the "Step-Function Pump" method. This is a dynamic technique, rather than a steady-state one.

First, we define a characteristic time for the cavity, t m t_{m}, which is the mean lifetime of a photon in the empty cavity. This is related to the cavity's Q-factor and losses, and is typically in the range of milliseconds to microseconds.

The experiment proceeds as follows. We start with the laser pump turned off. Then, at time t = 0 $t = 0$, we apply the pump power very rapidly, ideally as a perfect step-function.

At this moment, all of the possible laser modes within the gain profile begin to grow from the background of spontaneous emission noise. They start a race to build up their intensity, all competing for the same pool of gain.

Crucially, we do not wait for the laser to reach a steady state. Instead, we use a fast optical switch, like an acousto-optic modulator or AOM, to act as a gate. We open this gate for a very short time at a specific delay, t t, after the pump was turned on. This delay time t t is chosen to be less than the mode lifetime t m t_{m}.

This gated pulse of light, which represents a snapshot of the laser spectrum at time t t, is then sent to a spectrometer for analysis. For each mode q q in the spectrum, its intensity will depend on how quickly it was able to grow in that short time interval.

Page 58:

The growth of the power in each mode 'q', P q (t) $P_{\mathrm{q}}(t)$, follows a specific mathematical law. The equation is shown on the slide:

P q (t) = P q (0) , its initial noise power, × t π t m × (a large exponential term) .

$$P_{\mathrm{q}}(t) = P_{\mathrm{q}}(0),\ \text{its initial noise power, } \times \frac{t}{\pi\, t_{\mathrm{m}}} \times \text{(a large exponential term).}$$

Let's focus on the exponent, which contains two parts. The first part, involving Δ ω g $\Delta\omega_{\mathrm{g}}$, describes a process called spectral narrowing, where the overall laser spectrum becomes narrower as time goes on. The second part is the key to the sensitivity: it's e − α (ω q) c t $e^{-\alpha(\omega_{\mathrm{q}})\,c\,t}$.

This second term has the exact mathematical form of the Beer–Lambert law, I = I 0 e − α L $I = I_0\, e^{-\alpha L}$. By comparing these forms, we can see that in time-resolved ICLAS, the effective interaction length, L eff L_{eff}, is simply the speed of light, c c, multiplied by the generation time, t t.

L eff = c t .

$$L_{\mathrm{eff}} = c\, t.$$

This leads to some astounding numbers. Let's take the numerical example. Suppose we let the laser evolve for a generation time 't' of just 10 − 4 s $10^{-4}\,\mathrm{s}$, or 100 μ s $100\,\mu s$. The effective path length is L eff = c t $L_{\mathrm{eff}} = c\, t$, which is 3 × 10 8 m / s × 10 − 4 s $3 \times 10^{8}\,\mathrm{m/s} \times 10^{-4}\,\mathrm{s}$. This gives an effective path length of 3 × 10 4 m $3 \times 10^{4}\,\mathrm{m}$, or 30 k m $30\,\mathrm{km}$!

With a 30 k m $30\,\mathrm{km}$ effective path length, even a very weak absorption becomes easily detectable. If our detection system can measure a 1 % 1% dip in the intensity of a mode, this corresponds to a minimum detectable absorption coefficient, α min $\alpha_{\min}$, of approximately 3 × 10 − 9 c m − 1 $3 \times$

$10^{-9}\,\mathrm{cm}^{-1}$. This demonstrates the extraordinary sensitivity that can be achieved with this time-domain approach.

Page 59

This diagram provides an excellent visualization of the Time-Resolved ICLAS, or step-function pump, method.

The top graph illustrates the timing sequence. The horizontal axis is time. The blue line shows the pump power, which is off before t = 0 $t = 0$, and then instantly steps up to a constant value. Below this timeline, we see two red and orange boxes representing two different detection windows, one at an early time t 1 t_1, and another at a later time t 2 t_2.

The bottom graph shows the resulting laser spectrum at these two moments in time. The vertical axis is intensity and the horizontal axis is frequency. The red curve represents the laser spectrum as measured at the early time, t 1 t_1. At this point, the laser has just started to turn on. The spectrum is still relatively broad. We can see two small dips in the spectrum, which correspond to two absorption lines in our intracavity sample.

The orange curve represents the spectrum measured at the later time, t 2 t_2. Two important things have happened during the time interval between t 1 t_1 and t 2 t_2. First, you can see that the overall width of the laser spectrum has decreased. This is the "spectral narrowing" effect. Second, and most importantly, look at the absorption dips. They have become dramatically deeper and more pronounced. The effect of the absorption is amplified over time, as the non-absorbing modes have had more time to out-compete the absorbing modes. This diagram beautifully illustrates how the absorption contrast grows over the generation time, leading to the technique's high sensitivity.

Page 60:

Let's now consider a classic and very elegant practical example that showcases the power and selectivity of ICLAS: the detection and differentiation of iodine isotopes.

The experiment is set up as follows. We place an intracavity cell that contains a trace amount of the iodine isotope I 127 I_{127}, in its molecular form I 2 I_2. This cell is placed inside the cavity of a broadband dye laser.

This dye laser is designed to be multimode, with a gain profile so broad that its output spectrum simultaneously covers the absorption bands of both the stable isotope, I 127 I_{127}, and another isotope, the radioactive I 129 I_{129}.

Now, because of the powerful mode competition effect we discussed, the specific laser modes whose frequencies coincide with the absorption lines of the I 127 I_{127} that is inside the cavity will be strongly suppressed.

To visualize this effect, we now use the output of this modified laser as a probe beam. We send this beam to two *external* fluorescence cells. Cell A contains a pure sample of I 127 I_{127} vapor. Cell B contains a pure sample of I 129 I_{129} vapor.

What will happen in Cell A? The light entering Cell A from our ICLAS laser is specifically missing the frequencies required to excite I 127 I_{127} molecules. The pump lines have been extinguished by the intracavity absorption. Therefore, the I 127 I_{127} in Cell A cannot be excited, and we will observe no fluorescence.

Page 61:

Now, what happens in the second external cell, Cell B, which contains pure I-129 vapor?

The light coming from our ICLAS laser is only missing the frequencies corresponding to I-127 absorption. The laser is still oscillating strongly at all the frequencies corresponding to the absorption lines of the I-129 isotope. Therefore, when this light passes through Cell B, the I-129 molecules will readily absorb the light and will be excited, leading to the emission of strong fluorescence.

The final result is striking and demonstrates the incredible selectivity of the technique. We see strong fluorescence from Cell B, but no fluorescence from Cell A. This experiment demonstrates isotope-specific extinction. We have effectively created a filter that removes light at the frequencies of one isotope, while leaving the light at the frequencies of another isotope untouched.

The enhancement factor, Q, in such experiments can be enormous, on the order of 10 5 10^5. This highlights the ability of ICLAS to detect and differentiate species with extremely high sensitivity and specificity.

Page 62:

This diagram provides a clear visual summary of the iodine isotope selectivity experiment we've just described.

On the far left, we start with our "Broadband Dye Laser," which has the intracavity cell containing the I-127 sample. This produces a "Modified Laser Output."

The spectrum of this modified output is shown in the graph at the top. The overall shape is the broad gain profile of the dye laser. However, superimposed on this are sharp, dark lines. These are the "Suppressed modes at 127-I-2 frequencies." The intracavity absorption has effectively punched holes in the laser spectrum. Meanwhile, the "Modes still present at 129-I-2 frequencies" remain at full intensity.

This modified laser beam is then directed to the two "External Fluorescence Cells." The beam enters Cell A, which contains pure I-127. Since the pump lines for I-127 are missing from the beam, the result is "NO FLUORESCENCE." The beam then enters Cell B, which contains pure I-129. The pump lines for I-129 are present in the beam, so the result is "STRONG FLUORESCENCE."

This simple and elegant diagram perfectly captures the power of multimode ICLAS to act as an ultra-sensitive, species-specific filter.

Page 63:

Like any experimental technique, ICLAS has a unique set of advantages and limitations. Let's summarize the pros and cons.

First, the Pros:

The primary advantage is the enormous effective path length, Leff L_{eff}, that can be achieved. As we saw in the time-resolved case, this can easily reach tens or even thousands of kilometers, providing unparalleled sensitivity for detecting weak absorptions.

Second, when using a multimode laser, ICLAS allows for simultaneous broadband measurement. We can capture an entire spectral region in a single laser shot, which makes the technique very fast and efficient for spectroscopic surveys.

Third, because of its extreme sensitivity, ICLAS is an excellent tool for studying very weak spectral features that are inaccessible with other methods. This includes high overtone bands of molecules and quantum-mechanically forbidden transitions, which are often of great interest for fundamental physics and chemistry.

Now for the Cons:

The main drawback is that the technique's power is also its vulnerability. Because it relies on the delicate balance of the laser's internal dynamics, it requires a very stable laser cavity. The system is extremely sensitive to any optical feedback from outside the cavity and requires meticulous alignment and mechanical and thermal stability.

Page 64:

Continuing with the limitations of ICLAS:

Near-threshold operation, while providing the highest theoretical sensitivity, is also the regime where the laser is most prone to intensity noise and to "mode hops," where the laser frequency jumps uncontrollably between different longitudinal modes.

Furthermore, the sample itself can perturb the laser's operation in ways beyond simple absorption. A sample that is highly scattering, or one that causes significant gain depletion, can disrupt the lasing process itself, making the measurement unreliable.

A significant practical limitation is that the laser's gain medium and the sample cell must coexist within the same optical cavity. This often means they must share the same pressure and thermal environment, which can limit the range of conditions that can be studied. For example, it is difficult to perform ICLAS on high-pressure samples if the gain medium requires low pressure to operate.

Fortunately, there are mitigations for these problems. A major one is to move from an active ICLAS setup to a passive cavity setup. Instead of placing the sample inside the active laser cavity, we can use an external passive enhancement cavity, like the Fabry-Pérot systems we discussed earlier. We can then use an optical isolator between the laser and the passive cavity to completely prevent any destabilizing feedback from reaching the laser, giving us the best of both worlds: a stable laser and a high-power environment for our sample.

Page 65:

Here are a couple more advanced mitigation strategies to overcome the limitations of ICLAS.

One issue with standard linear laser cavities is a phenomenon called "spatial hole burning." The counter-propagating beams inside a linear cavity create a standing wave pattern of the electric field, which has nodes (points of zero intensity) and antinodes (points of maximum intensity). If our absorbing atoms or molecules happen to be located at the nodes, they will not interact with the laser field, reducing the sensitivity. A very effective way to avoid this is to employ a ring-laser geometry. In a ring laser, the light is constrained to travel in only one direction around a closed loop. This creates a traveling wave, not a standing wave, which has a uniform intensity profile. This ensures that all the absorbers within the beam path interact equally with the laser field.

Another mitigation, which is essential for any high-performance system, is to implement active stabilization. Instead of just trying to build a passively stable system, we can use electronic feedback loops to actively control the critical parameters. This includes stabilizing the power of the pump source and actively stabilizing the length of the laser cavity using the very locking techniques we've already discussed. By actively fighting against sources of noise and drift, we can achieve robust, high-sensitivity operation.

Page 66:

This slide provides an excellent summary of the enhancement factors for the various techniques we have discussed, allowing us to compare their power and effectiveness on a common scale.

Let's start with our baseline: a simple, single-pass absorption measurement. Here, the effective path length, Leff L_{eff}, is just the physical length, L L. We define its enhancement factor, Q Q, as being equal to 1 1.

Next, the multipass cell. Here, L e f f L_{eff} is q L qL, where q q, the number of passes, is typically in the range of 50 to 200. This is a simple, robust, and significant improvement.

Third, the external high-Q cavity. Here, the power gain inside the cavity is approximately 1 1 − R $\frac{1}{1-R}$. The effective path length enhancement is on the order of 2 R 1 − R $\frac{2R}{1-R}$. For a typical high-quality mirror, this also results in an enhancement factor of a few hundred, similar to a multipass cell, but with the added benefit of a massive power buildup.

Fourth, FM derivative detection. This enhances sensitivity in a different way. It doesn't increase the path length, but it reduces the noise bandwidth by using a lock-in amplifier. This results in an effective signal-to-noise ratio, or SNR, improvement of typically 10 to 1000 times, depending on the modulation frequency.

Finally, intracavity techniques. For a single-mode ICLAS system, the enhancement factor Q Q is equal to 1 T 2 $\frac{1}{T_2}$, the output coupling, which is typically in the range of 50 to 500. This can be even higher when operated near the laser's threshold.

Page 67:

Continuing our summary of enhancement factors, we now come to the most powerful techniques.

For an intracavity system that is multimode and strongly coupled, we start with the basic intracavity enhancement factor, but this is then multiplied by the additional mode competition factor, which is 1 + K N $1 + KN$. This can provide an additional multiplicative factor of 10 3 10^3 to 10 4 10^4. This combination leads to truly enormous overall enhancement factors.

Finally, for time-resolved intracavity spectroscopy, the enhancement comes from the effective path length being equal to c t ct. This can result in a path length enhancement of 10 5 10^5 to 10 8 10^8 times that of a single pass. These numbers are staggering and represent the pinnacle of path-length enhancement techniques.

This summary clearly shows the hierarchy of techniques, from simple multipass cells providing modest gains to advanced intracavity methods that push sensitivities into truly extraordinary regimes.

Page 68:

This slide transitions us from theory to the practical reality of building an ultra-sensitive spectroscopic setup. It provides an excellent and comprehensive design checklist that anyone embarking on such a project should consider.

Let's go through the main categories:

First, the Source. You need a high-quality laser that is single-mode, has a narrow linewidth, and is tunable over the spectral range of interest. Critically, you must use an output isolator to prevent any back-reflections from destabilizing the laser.

Second, the Optics. For any cavity-based method, precise mode-matching is absolutely essential. You should aim for a coupling efficiency of 95% or better. And your mirror coatings must be superb, with very high reflectivity and, just as importantly, extremely low scattering loss, as scattered light can create spurious interference fringes.

Third, the Electronics. If using an EOM, you need a high-frequency driver with low phase noise. Your detection chain—the photodiode and the transimpedance amplifier—must be low-noise and have sufficient bandwidth for your modulation frequency. And a modern, digital lock-in amplifier with a selectable time constant is an indispensable tool.

Fourth, Sample Handling. You need a well-designed sample cell, perhaps one that is pressure-tunable and has temperature control. You also need to be meticulous about gas purification to remove any parasitic absorbers, like water vapor, that could create interfering signals.

Finally, Data Processing. Your job isn't done when you acquire the data. You need robust software to perform a simultaneous fit of your data, perhaps using multiple derivative orders, to accurately extract the absorption coefficient α (ω) $\alpha(\omega)$. And crucially, you must always perform a calibration of your system using a known reference gas, such as acetylene, which has a well-characterized and strong absorption spectrum.

Page 69:

To bring all these concepts together, let’s walk through an example experimental workflow for a state-of-the-art experiment, such as a Noise-Immune Cavity-Enhanced FM spectroscopy setup.

Step 1: Alignment. This is often the most painstaking part. First, you must align your laser beam to pass cleanly through your electro-optic modulator, or EOM. Then, it goes through the mode-matching telescope. Finally, you must align this shaped beam to perfectly couple into your high-finesse enhancement cavity, which contains your absorption cell.

Step 2: Lock the cavity. Once you have light successfully building up inside the cavity, you need to engage the feedback lock to keep it on resonance. You would typically use the Pound-Drever-Hall, or PDH, technique. This method cleverly leverages the existing FM sidebands that are generated by your EOM to create a very robust error signal, so this step integrates naturally into an FM experiment.

Step 3: Set the modulation index. With the cavity locked, you now need to optimize the modulation parameters. You would adjust the radio-frequency power going to your EOM to set the modulation index, a a, to an optimal value. A common strategy is to set it such that the power in the carrier (related to the Bessel function J 0 (a) $J_0(a)$) is approximately equal to the power in the first-order sidebands (related to J 1 (a) $J_1(a)$). This provides a good balance for generating a strong signal.

Page 70: Continuing with our experimental workflow

Step 4: Sweep and record. With the system aligned, locked, and optimized, you are now ready to acquire data. You would program your laser to sweep its center frequency slowly across the target molecular transition. While the laser is sweeping, you record the output signal from your lock-in amplifier. This will trace out the characteristic derivative-shaped spectrum of your absorption line.

Step 5: Convert the signal. The raw data you record is a voltage from the lock-in amplifier as a function of laser frequency. The next crucial step is to convert this derivative signal back into an absolute absorption spectrum, α(ω) $\alpha(\omega)$. This requires a calibrated physical model of your experiment. You need to know your modulation index a a, the effective path length Leff L_{eff} provided by your cavity, and the responsivity of your detector and electronics. By fitting your derivative data to a theoretical model, you can extract the quantitative absorption coefficient.

Step 6: Validate. This is the final and most important step for any quantitative measurement. You must validate your entire system and your analysis model. This is typically done by injecting a calibrated gas mixture with a precisely known concentration into your cell. You then perform the measurement and check if your result matches the known value. If there is a discrepancy, you can refine your model and your noise-equivalent power (NEP) budget until your system provides accurate and reliable results.

Page 71:

As we draw this chapter to a close, let's reflect on some concluding remarks.

The first, and perhaps most important, takeaway is that high-sensitivity absorption spectroscopy is not a single discipline. It is a deeply multi-disciplinary optimisation problem. To design, build, and operate a state-of-the-art experiment, one must be an expert in three distinct but interconnected fields:

1. Photonics: This includes the design and alignment of high-Q optical cavities, understanding Gaussian beam optics for mode-matching, and selecting and characterizing low-noise, narrow-linewidth lasers.

2. Electronics: This involves the world of radio-frequency and microwave engineering for generating fast and stable modulation, designing low-noise detection circuits, and mastering the theory and practice of phase-sensitive lock-in detection.

3. Spectroscopic theory: This is the physics of interpreting the data. It requires a deep understanding of lineshape theory, the ability to model the complex derivative signals generated by modulation techniques, and the statistical methods required to fit experimental data to these models to extract meaningful, quantitative results.

It is the synergy and mastery of all three of these areas that allows experimentalists to push the frontiers of what is detectable.

Page 72:

The second major conclusion is that the powerful combination of our two main strategies—path-length enhancement and noise-rejection modulation—routinely achieves detection limits that were previously thought to be impossible. These are not just incremental improvements; they represent many orders of magnitude in progress, allowing us to move from detecting parts-per-million to parts-per-billion, parts-per-trillion, and in some cases, even parts-per-quadrillion.

Mastery of these techniques is far more than just an academic exercise in pushing instrumental limits. It opens the doors to transformative research and new capabilities across a vast range of scientific and technological fields.

For example:

1. In climate science, these methods are used for ultra-trace detection of greenhouse gases, atmospheric pollutants, and reactive chemical intermediates that govern the chemistry of our atmosphere.

2. In fundamental physics, they are used to perform some of the most precise tests of our physical laws, by searching for tiny, parity-violating energy shifts in chiral molecules, or by measuring the frequencies of forbidden transitions in simple atoms to test the predictions of quantum electrodynamics.

3. In medical diagnostics, there is a burgeoning field of breath analysis, where the trace volatile organic compounds in a patient's breath are measured to provide a non-invasive signature of metabolic processes and diseases.

Page 73:

The applications of these powerful techniques extend even further.

4. In engineering, they are used for combustion and plasma monitoring. These methods allow us to peer inside the harsh environment of an internal combustion engine or an industrial plasma reactor to monitor chemical reactions and optimize efficiency in real-time.

Finally, let's look to the future. This field is by no means static. Continual technological advances promise to deliver even greater sensitivity and functionality in the coming years.

For example, the development of optical frequency combs provides a way to perform massively parallel, broadband spectroscopy with thousands of ultra-precise laser lines at once. The development of whispering-gallery-mode micro-cavities allows us to create resonators with astronomical Q-factors on the scale of a microchip. And the field of integrated photonics promises to move these complex, table-sized experiments onto a single, compact, and robust photonic chip.

These advances promise a new generation of ultra-sensitive spectrometers that will find applications we haven't even dreamed of yet. The quest to measure the smallest signals and to detect the undetectable is a journey that continues to be at the very heart of experimental physics.

Thank you.