Vol. 2

Chapter 1.2.4

Page 1:

Good morning, everyone. Welcome to this section of our graduate course in Laser Spectroscopy. I'm Distinguished Professor Dr M A Gondal, and it's a pleasure to have you in this advanced graduate physics course.

Today, we embark on a fascinating and incredibly powerful topic, one that represents a significant leap forward in the art and science of measuring vanishingly small absorptions. As you can see from the title slide, we will be delving into Chapter 1.2.4 of our course notes, focusing on a technique that has truly revolutionized trace species detection.

Page 2:

The topic for today's lecture is Cavity Ring-Down Spectroscopy, which is almost universally known by its acronym, C. R. D. S.

Now, before we dive into the mathematics and the experimental schematics, I want you to hold a central idea in your mind.

The genius of CRDS lies in a fundamental shift in perspective. For over a century, traditional absorption spectroscopy, governed by the Beer-Lambert law, has been about measuring a *change in light intensity*. You measure how much light you start with, how much light gets through your sample, and the difference tells you about the absorption. This is simple and effective, but it has a critical weakness: it is exquisitely sensitive to any fluctuations in your light source. If your laser power flickers, your measurement is compromised.

CRDS sidesteps this problem entirely. It doesn't measure intensity; it measures *time*. Specifically, it measures the characteristic time it takes for light to die out, or "ring down," inside a highly reflective optical cavity. As we will see, this time constant is directly related to the total losses within the cavity, including the absorption from our sample of interest. By measuring a time constant instead of an intensity ratio, CRDS becomes almost completely insensitive to the very laser power fluctuations that plague traditional methods. This is the secret to its extraordinary sensitivity and precision. So, let's begin our detailed exploration of how this remarkable technique works.

Page 3:

Alright, let's formally define what we're talking about. What *is* Cavity Ring-Down Spectroscopy?

As the first bullet point states, CRDS is an absorption-measurement technique. Its primary application, and where it truly excels, is in determining the concentration of *very weakly absorbing species*. Think about trying to detect parts-per-billion or even parts-per-trillion of a pollutant in the atmosphere, or measuring a highly forbidden molecular transition. These are signals that would be completely buried in the noise of a conventional spectrometer.

The core of the method is to monitor how rapidly light energy "rings down"—or, more formally, decays—inside a what we call a high-finesse optical cavity. We'll define finesse more rigorously later, but for now, just think of a cavity made of two of the best mirrors money can buy, mirrors that are so reflective that a photon of light can bounce back and forth between them thousands, or even tens of thousands, of times before it's lost.

This brings us to the second, crucial point, which I alluded to in my introduction. Instead of measuring an absolute transmitted intensity, CRDS measures a *time constant*, which we'll denote with the Greek letter τ.

Think about it like striking a bell. A high-quality bell rings for a long time; its sound decays slowly. If you put your hand on the bell, you introduce a new loss mechanism—damping—and the sound dies out much more quickly. The bell's "ring-down time" gets shorter. In CRDS, our optical cavity is the bell, the light is the sound, and the absorbing gas sample inside the cavity is the hand that damps the ringing. By measuring how much the ring-down time shortens, we can quantify the absorption. The asterisk here highlights the key benefit: this method inherently removes most errors caused by laser power fluctuations. The initial brightness of the light pulse doesn't affect the *rate* at which it decays, just as how hard you strike the bell doesn't affect the pitch or the decay rate, only the initial volume.

Page 4

Let's continue building this physical picture. How is a measurement actually performed?

First, a short laser pulse is injected into the space between two mirrors. The defining characteristic of these mirrors is their incredibly high reflectivity, which we denote with a capital R. As the slide notes, this value is extremely close to unity—typically, $R \geq 0.999$. That's 99.9% reflective, and often we use mirrors that are 99.99% or even 99.999% reflective.

Because the mirrors are not perfectly reflective, a tiny, tiny fraction of the light escapes, or "leaks out," through the second mirror on each and every round-trip the pulse makes inside the cavity. A photodetector placed behind this second mirror will therefore see a train of successively weaker pulses, one for each time the main pulse hits that mirror. It is the exponential decay of the intensity of this leakage train that we record and analyze.

Now, consider the second bullet point. This is the heart of the measurement. In an empty, pristine cavity, the decay time is determined solely by the intrinsic losses of the mirrors themselves—their transmission and any scatter or absorption losses. We can measure this "empty cavity" decay time, let's call it τ_0. Then, we introduce our sample—a gas, for instance—into the cavity. This gas provides a new channel for energy loss: molecular absorption. Any additional absorption inside the cavity provides another damping mechanism, which, as the slide says, *slightly shortens the decay time*. We measure this new, shorter decay time, let's call it τ.

By simply comparing these two decay times—the one with the absorber and the one without—we can isolate the loss due solely to the absorber. This comparison directly yields the sample's absorption coefficient, which we denote with the Greek letter α. The entire measurement boils down to two timing measurements.

Page 5:

This diagram provides a clear, simple schematic of the fundamental CRDS setup. Let's walk through it together.

On the far left, we have the source of a laser pulse, represented by the thick red line entering the system. This pulse travels towards and enters an optical cavity.

The cavity itself is defined by two concave mirrors, labeled Mirror 1 and Mirror 2, facing each other. The physical distance between the mirrors is the cavity length, labeled with a capital L L. You'll notice the specifications for the mirrors: R ≥ 0.999 $R \geq 0.999$. These are the super-mirrors we just discussed.

The laser pulse enters through Mirror 1 and becomes trapped inside the cavity. It then begins to bounce back and forth between the two mirrors. The blue curved arrows illustrate this bouncing path. The total distance traveled for one complete back-and-forth trip is, of course, twice the cavity length, so the round-trip distance is 2 L $2L$.

Each time the trapped pulse strikes Mirror 2, a tiny fraction of its energy leaks through. This is shown as the “Leakage Path.” This leaking light is captured by a Photodetector placed just behind Mirror 2.

So, in a single event, we inject a pulse. It bounces back and forth many, many times. And with each bounce on Mirror 2, a small pulse of light emerges, gets detected, and the photodetector records a signal that decays over time. It is this decay curve whose time constant we need to measure. The concept is beautifully simple, yet profoundly powerful.

Page 6:

Let's now consider the historical motivation for developing CRDS and get a quantitative feel for its sensitivity.

Traditional single-pass absorption measurements, where you simply shine a laser through a sample cell of length L L, run into trouble very quickly. As the first bullet point says, they struggle when the absorption length, capital L L, is limited—perhaps you only have a small sample cell—or, more fundamentally, when the absorption coefficient, α α, is so small that the total absorption, which is the product α L αL, is smaller than the noise on your detector. If your laser power fluctuates by one percent, you can't possibly hope to measure an absorption of zero point one percent.

To combat this, researchers developed early "long-path" methods. You may have heard of Herriott cells or White cells. These are clever optical systems using specially shaped mirrors to fold a long optical path into a compact physical volume. They effectively increase the interaction length L L by making the beam pass through the sample many times—perhaps 50 or 100 times. This increases the total absorption, α L αL, making it easier to measure. However, these multipass cells can be very tricky to align and are often sensitive to vibrations.

This is where CRDS entered the scene, primarily from the mid-1990s onward, and raised the bar for sensitivity by orders of magnitude. The reason is stated in the third bullet point, and it's a critical concept. The *effective path length*, L e f f L_{eff}, in CRDS is given by the physical length of the cavity, L L, divided by the total losses per pass, which for an empty cavity is approximately one minus the reflectivity, R R. So, L e f f L_{eff} equals L L divided by the quantity one minus R R.

L e f f = L 1 − R .

$$L_{\text{eff}} = \frac{L}{1 - R}.$$

Let's put some numbers to this.

Page 7:

Let's continue that thought. If we have a cavity that is just one meter long—a very manageable size for a lab bench—and we use mirrors with a reflectivity R R of 0.9999 0.9999 (that's four nines), what is our effective path length?

Well, 1 − R = 1 − 0.9999 $1 - R = 1 - 0.9999$, which is 10 − 4 10^{-4}. So, L e f f L_{eff} is one meter divided by 10 − 4 10^{-4}. This gives an effective path length of 10 4 10^{4} meters, which is ten kilometers! We have created a ten-kilometer-long absorption cell on a one-meter optical bench. This is the magic of the high-finesse optical cavity. It effectively folds the optical path thousands of times.

This enormous path length amplification means even a tiny absorption coefficient, α α, results in a measurable effect. Furthermore, as we've stressed, the measurement itself is a time constant. And time constants can be measured with sub-percent precision, even when the signal itself is very weak. A decaying exponential is a mathematically simple and robust function to fit.

The result of all this is truly astonishing achievements in sensitivity. The last bullet point notes that detection limits for the absorption coefficient, α α, can get down to approximately α ≈ 10 − 11 c m − 1 / H z .

$$\alpha \approx 10^{-11}\,\text{cm}^{-1}/\sqrt{\text{Hz}}.$$

Let's unpack that unit. Inverse centimeters is the standard unit for α α. The "per root Hertz" part tells us about the averaging time. It means that with one second of signal averaging, we can achieve a sensitivity of 10 − 11 10^{-11} inverse centimeters. This allows for the detection of extremely rare molecules or very weak transitions that were previously inaccessible.

Page 8:

This bar chart provides a wonderful visual summary of the evolution of spectroscopic sensitivity, putting CRDS into context. Let's analyze it carefully.

The horizontal axis represents different spectroscopic techniques. The vertical axis is the Minimum Detectable Absorption, which we call α min $\alpha_{\min}$, plotted on a logarithmic scale in units of c m − 1 cm^{-1}. Remember, a lower value on this chart means higher sensitivity, which is what we're striving for.

Let's start on the left. The first bar is for a "Single-Pass" measurement. Here, the effective path length, L e f f L_{eff}, is just the physical length of our sample cell, say ten centimeters. The minimum detectable α α is around 10 − 4 10^{-4}. This is respectable, but not great for trace detection.

Moving to the next bar, "Multipass." This represents techniques like the Herriott and White cells we mentioned. By folding the path, we increase L e f f L_{eff} to somewhere between 10 and 100 meters. As a direct result, our sensitivity improves dramatically. The α min $\alpha_{\min}$ drops by about two orders of magnitude, down to around 10 − 6 10^{-6}. A significant improvement.

The third bar is for "Intracavity Laser Absorption Spectroscopy," or ICLAS. This is another very sensitive technique where the absorbing sample is placed inside the laser cavity itself. It relies on different principles involving mode competition, but it can achieve very high equivalent path lengths, noted here as greater than one kilometer. The sensitivity is correspondingly better, with α min $\alpha_{\min}$ pushing down towards 10 − 8 10^{-8} or 10 − 9 10^{-9}.

And finally, on the far right, we have our technique of the day: CRDS. Look at the immense improvement. The effective path length is now on the order of ten kilometers, as we calculated. This pushes the minimum detectable absorption down to 10 − 10 10^{-10}, 10 − 11 10^{-11}, and even approaching 10 − 12 10^{-12}. Each step on this chart represents a massive leap in our ability to see the invisible, and CRDS currently stands at the pinnacle for this type of absorption measurement.

Page 9:

Alright, let's now transition from the conceptual overview to the mathematical formalism. To do that, we need to precisely define the basic elements and quantities associated with an optical ring-down cavity.

First, as we've seen in the diagrams, the core of the system is two mirrors. We'll assume for simplicity that they are identical, with reflectivity R 1 = R 2 $R_1 = R_2$, which we'll just call R R. These mirrors face each other to form a linear cavity of physical length L L.

Now, for any real mirror, reflectivity isn't the whole story. The energy of a photon incident on the mirror has to go somewhere. The first bullet point under "Other characteristic quantities" notes the transmission of each mirror, T T. But there are other loss mechanisms. So, we have a fundamental conservation of energy relationship for the mirror, which is the equation shown:

T = 1 − R − A

$$T = 1 - R - A$$

Let's break this down. One represents one hundred percent of the incident energy. R R is the fraction that is reflected. T T is the fraction that is transmitted through the mirror. And capital A A represents the fraction of energy that is lost to all other processes *at the mirror surface*.

Page 10:

Continuing with that thought, what does this loss term, capital A, physically represent? As the first line on this slide clarifies, capital A is a catch-all term that lumps together all non-sample, non-transmission losses associated with the mirror. This includes light that is *scattered* off the mirror surface due to microscopic imperfections, light that is lost to *diffraction* because the mirror has a finite size, and any light that is directly *absorbed* by the mirror's dielectric coating itself. For the ultra-high-quality mirrors used in CRDS, this 'A' term is incredibly small, but it's not zero, and it ultimately sets a limit on the performance of the empty cavity.

Next, we define two fundamental temporal and spatial parameters. The round-trip optical path is simply 2 L $2L$, as we saw in the diagram. From this, we get the round-trip time, denoted T R T_R. This is the time it takes for a photon to travel from one mirror, to the other, and back again. The equation is simple: T R T_R equals the round-trip distance, 2 L $2L$, divided by the speed of light, c c. So, T R T_R equals 2 L c $\frac{2L}{c}$. For a one-meter cavity, the round-trip time is about 6.7 nanoseconds. This is a very important timescale in our analysis.

Finally, we must state a crucial assumption that underpins all the simple derivations that follow. We assume that the mirrors are near-perfect. This means that the fraction of light transmitted, T T, and the fraction lost to scatter and absorption, A A, are both very, very much less than one. This is equivalent to saying that the reflectivity, R R, is extremely close to one, which is the entire premise of CRDS.

Page 11

Now for a slightly more subtle point regarding the nature of the laser excitation.

The first bullet point here states that a short input pulse of initial power P 0 P_0 excites many cavity modes simultaneously if its bandwidth exceeds the cavity's free spectral range. Let's unpack this statement, because it's dense with important concepts.

First, what is a cavity mode? A stable optical cavity will only support standing waves of light at specific resonant frequencies, much like a guitar string only vibrates at its fundamental frequency and its harmonics. These allowed frequencies are the “modes” of the cavity. For a simple linear cavity, these modes are separated in frequency by a value called the Free Spectral Range, or FSR, which is equal to c 2 L $\frac{c}{2L}$.

Second, what is the bandwidth of a short laser pulse? From the uncertainty principle, or more formally, the Fourier transform relationship between time and frequency, a pulse that is short in time must be broad in frequency. A picosecond or femtosecond laser pulse has a very wide spectral bandwidth.

So, if the laser pulse’s bandwidth is wider than the spacing between the cavity modes (the FSR), the pulse has frequency components that will match and excite many of these cavity modes at the same time. While this is a very common scenario in pulsed CRDS, it can lead to complications like interference beats, which we will discuss later.

The slide also foreshadows an alternative approach: mode-matching to a single T E M 00 TEM_{00} mode. T E M 00 TEM_{00} refers to the fundamental transverse Gaussian mode, the “cleanest” spatial profile a laser beam can have. In some advanced CRDS schemes, you use a very narrow-bandwidth laser and carefully match its beam profile and frequency to just one of these cavity modes. For now, we will proceed with the simpler picture of a short, broadband pulse.

Page 12:

Let's begin our journey of deriving the mathematical expression for the ring-down signal. We'll start with the very first pulse of light that leaks out of the cavity. We'll call its power P 1 P_1.

The first bullet point reminds us that immediately after the initial laser pulse is injected into the cavity, a part of it exits through the output mirror. How much power does it have?

To figure this out, we first need to define the absorption coefficient of our sample inside the cavity. This is given by the Greek letter α α, typically in units of inverse centimeters, or centimeters to the minus one. This α α is the quantity we ultimately want to measure.

Now, let's consider the power in this very first transmitted pulse, P 1 P_1. The equation is given as:

P 1 = T 2 e − α L P 0

$$P_1 = T^2 e^{-\alpha L} P_0$$

Let's deconstruct this piece by piece.

Page 13:

Let's break down that equation for the power of the first pulse, P 1 P_1.

The first term is T 2 T^2. Where does this come from? The initial laser pulse, with power P 0 P_0, must first *enter* the cavity. It does so by passing through the first mirror, so its power is multiplied by the mirror's transmission, T T. Then, this pulse travels a single pass down the length of the cavity, L. It then strikes the second mirror and *exits* the cavity, passing through it. This involves another multiplication by the transmission, T. So, the T^2 term represents the product of two transmission events: one on entry, and one on exit.

The next term is the exponential: $e^{-\alpha L}$. This is nothing more than the familiar Beer–Lambert attenuation law. It describes the fraction of light that survives after travelling a distance L through a medium with an absorption coefficient α. This accounts for the absorption by our sample during that first single pass.

The final term, P_0, is simply the incident laser power of the pulse we started with.

So, to reiterate, the power of the very first leakage pulse, P_1, is what you get after the initial pulse enters the cavity (factor T), traverses the sample once (factor $e^{-\alpha L}$), and immediately exits (factor T). This gives us our starting point for the subsequent decay.

Page 14:

Now that we have the power of the first leaked pulse, P_1, what happens to the light that remains trapped inside the cavity? Let's consider its evolution after n full round-trips.

The first bullet point breaks down the losses that occur during a single round-trip. After entering the cavity, the pulse travels to the far mirror and back. This involves two mirror reflections. Since the reflectivity of each mirror is R, the power is multiplied by R^2 after one round-trip. Simultaneously, the pulse has traversed the sample twice—once down the length L L, and once back. So it travels a total distance of 2 L $2L$ through the absorbing medium. The attenuation factor for this is therefore the exponential of minus two alpha L.

So, for each complete round-trip, the intra-cavity power is multiplied by a total factor of (R 2 e − 2 α L) $(R^2e^{-2\alpha L})$.

Therefore, as the second bullet point states, just before the (n + 1) $(n+1)$-th leakage event—which means after n n full round-trips have been completed since the first leakage—the power of the pulse still inside the cavity has been multiplied by this factor n n times. That is, it's been multiplied by the quantity (R 2 e − 2 α L) $(R^2e^{-2\alpha L})$ all raised to the power of n n.

This leads directly to the expression for the transmitted power of the n-th leakage pulse, which we denote P n P_n. It is simply the power of the first leakage pulse, P 1 P_1, multiplied by this round-trip loss factor n n times.

So we have:

P n = (R 2 e − 2 α L) n P 1

$$P_\mathrm{n} = (R^2e^{-2\alpha L})^n \, P_1$$

This expression is correct, but it's a bit cumbersome. To get it into the nice exponential decay form we're looking for, it's very convenient to rewrite it using logarithms, which we will do on the next slide.

Page 15:

Let's now perform the mathematical manipulation to transform our expression for P n P_n into a more useful exponential form.

First, we take the natural logarithm of both sides of the equation from the previous slide. The natural log of P_n, written $\ln P_n$, is equal to the natural log of P_1 plus the natural log of the term in parentheses. Using the properties of logarithms, the exponent 'n' comes down in front. So, we get:

ln P n = ln P 1 + n (2 ln R − 2 α L) .

$$\ln P_n = \ln P_1 + n\,(2\ln R - 2\alpha L).$$

Next, we can re-exponentiate this expression to write P_n in terms of the exponential function. P_n is equal to P_1 times the exponential of the second term. So:

P n = P 1 exp [2 n (ln R − α L)] .

$$P_n = P_1 \exp[2\,n(\ln R - \alpha L)].$$

Now comes the most important step in this derivation. We use the fact that the reflectivity, R, is very, very close to 1. This allows us to use a first-order Taylor series expansion for the natural logarithm. The expansion for $\ln(x)$ around $x = 1$ is $\ln(x)$ is approximately $x - 1$. So,

ln (R) ≈ R − 1.

$$\ln(R) \approx R - 1.$$

But we know from our earlier definition that the total losses at a mirror are given by $1 - R = T + A$. Therefore, $R - 1$ is simply the negative of $(T + A)$. So, we arrive at the crucial approximation:

ln (R) ≈ − (T + A) .

$$\ln(R) \approx -(T + A).$$

This simple linear approximation is the key that unlocks the final, elegant form of the decay expression.

Page 16:

Now we can take the result of that approximation and substitute it back into our equation for P_n.

The title of this slide says it all: we are expressing the "Power Decay in Terms of Losses $T + A + \alpha L$". We replace the term ‘$\ln R$’ with ‘$-(T + A)$’.

Our expression for P_n now becomes:

$$P_n = P_1 \exp[-2\, n\, (T + A + \alpha L)].$$

This equation has a wonderfully clear physical interpretation, which is stated in the second bullet point. The power decays exponentially with the number of round-trips, ‘n’. The argument of the exponential tells us the total fractional loss per round trip. This total loss is simply the sum of the individual loss mechanisms. For each round trip, we have a loss due to mirror transmission (which is $2T$, one for each mirror, but our formula normalizes this to a loss per pass length L, hence the structure), a loss due to mirror scatter and absorption (A), and a loss due to absorption by the medium (αL).

The beauty of this form is that all the loss terms now appear linearly and additively inside the exponential. This is exactly what we need to be able to separate them later.

Page 17:

This slide makes a final, crucial connection between our discrete pulse model and what is actually observed in many experiments.

The key insight is that the pulse number, ‘n’, is directly proportional to the observation time, ‘t’. We will formalize this on the very next slide, but the relationship is simple:

t = n T R

$$t = n\,T_R$$

Because of this direct proportionality, our equation for the power of the *n*-th pulse, P n P_n, can be rewritten as an equation for the power as a function of continuous time, P (t) $P(t)$.

Now, consider a typical experiment. The round-trip time, T R T_R, for a meter-long cavity is just a few nanoseconds. Many photodetectors and data acquisition systems have a response time that is slower than this. They cannot resolve the individual leakage pulses. What they see instead is the smoothed-out envelope of the pulse train.

Therefore, when detected with insufficient temporal resolution, the train of discrete leakage pulses merges into what appears to be a single, smooth exponential decay. And the time constant of this smooth decay is precisely what our equation describes.

Page 18:

Let's formalize the relationship between the discrete pulse number, n n, and the continuous physical time, t t. This is a straightforward but essential step.

First, let's recall the definition of the round-trip time, T R T_{R}. This is the time it takes for the light pulse to complete one full circuit of the cavity. It is given by the equation:

T R = 2 L c

$$T_{\mathrm{R}} = \frac{2\,L}{c}$$

Now, let's think about the arrival time of the n n-th leakage pulse at the detector. We'll set t = 0 $t = 0$ as the arrival time of the first leakage pulse. The second pulse arrives after one round-trip, at time T R T_{R}. The third arrives after two round-trips, at time 2 T R $2\,T_{\mathrm{R}}$. And so, the n n-th leakage pulse, which has completed n − 1 $n - 1$ additional round-trips, arrives at a time we can simply call 't'.

So, the arrival time 't' of the n n-th leakage pulse is given by:

t = n T R

$$t = n\,T_{\mathrm{R}}$$

Substituting our expression for T R T_{R}, we get:

t = n 2 L c

$$t = n\,\frac{2\,L}{c}$$

This simple linear relationship, t = n T R $t = n\,T_{\mathrm{R}}$, allows us to switch from the discrete variable 'n' to the continuous variable 't' in our decay equation.

Page 19:

Now we can put all the pieces together to arrive at the final expression for the measured signal.

The first bullet point summarizes the situation: If the detector response time, which we can call τ d e t τ_{det}, is much longer than the cavity round-trip time, T R T_{R}, then the individual pulses blur together into a smooth decay.

By replacing the term n n in our previous exponential equation with t t divided by T R T_{R}, which is t t divided by 2 L c $\frac{2L}{c}$, we find that the power as a function of time, P (t) $P(t)$, follows a simple exponential decay:

P (t) = P 1 exp (− t τ 1)

$$P(t) = P_1 \exp\left(-\frac{t}{\tau_1}\right)$$

This is the classic exponential decay curve that is fitted in a CRDS experiment. P 1 P_1 is the initial power at time t = 0 $t = 0$, and the crucial parameter is τ 1 τ_1.

What is τ 1 τ_1? As the slide states, τ 1 τ_1 is the "cavity ring-down time" when the absorber is present. It is the characteristic one-over-e time of the decay. Our entire goal is to measure this τ 1 τ_1, and its counterpart without the absorber, τ 2 τ_2, as accurately as possible.

On the next slide, we will derive an explicit expression for τ 1 τ_1 based on the physical parameters of the cavity.

Page 20:

This slide shows how we derive the explicit formula for the ring-down time, τ 1 τ_1. It's simply a matter of comparing the exponents from our last two expressions for the power decay.

In the previous slide, we wrote the decay term as exp (− t τ 1) $\exp(-t/\tau_1)$. Two slides ago, we derived the decay term as exp [− 2 n (T + A + α L)] $\exp[-2\,n\,(T + A + \alpha L)]$. If we substitute n = t T R = t c 2 L $n = t/T_R = t\,c/2\,L$ into the second expression, the decay term becomes exp [− t c L (T + A + α L)] $\exp[-t\ c/L\ (T + A + \alpha L)]$.

Now, we simply identify the coefficients of 't' in the exponents. This gives us the first equation on the slide:

1 τ 1 = c L (T + A + α L) .

$$\frac{1}{\tau_1} = \frac{c}{L}\,(T + A + \alpha L).$$

This equation tells us that the decay *rate* (1 / τ 1 $1/\tau_1$) is directly proportional to the sum of all the fractional losses per pass.

By simply rearranging this equation, we can solve for the ring-down *time*, τ 1 τ_1:

τ 1 = L / c T + A + α L .

$$\tau_1 = \frac{L/c}{T + A + \alpha L}.$$

This form is very intuitive. The numerator, L / c L/c, is the time it takes for light to make a single pass of the cavity. The denominator is the total fractional loss per pass. So the ring-down time is essentially the single-pass time scaled by one over the fractional loss.

Finally, we consider the very important special case: the empty cavity. In this situation, the absorption coefficient, α α, is equal to zero. We'll call the ring-down time for the empty cavity $\tau 2$ τ_2.

Page 21:

Continuing with the special case of the empty cavity, we simply set $\alpha = 0$ $\alpha = 0$ in our general expression for the ring-down time. This gives us the ring-down time for the reference measurement, which we've called $\tau 2$ τ_2.

The equation is:

$\tau 2 = L / c T + A$

$$\tau_2 = \frac{L/c}{T + A}$$

Here, the only losses are the intrinsic mirror losses: transmission T T and scatter/absorption A A.

Now, the slide shows a further approximation. If we assume that all the mirror losses, A A, are negligible compared to the transmission, T T, we can use the relation $T \approx 1 - R$ $T \approx 1 - R$. This gives the commonly cited simplified formula for the empty cavity ring-down time:

$\tau 2 \approx L / c 1 - R$

$$\tau_2 \approx \frac{L/c}{1 - R}$$

This form makes the dependence on mirror reflectivity very explicit. As $R \to 1$ $R \to 1$, the denominator approaches zero, and the ring-down time, τ τ, goes to infinity.

The remark at the bottom simply reiterates the basis for this approximation: we are using T + A = 1 − R $T + A = 1 - R$, which is valid when all other unaccounted-for losses are small. For our purposes, the more general form including A A is more accurate, and as we will see, the A A term will conveniently cancel out anyway.

Page 22:

We have finally reached the climax of our derivation. This is the slide where we see the true elegance and power of the CRDS method. Our goal is to extract the absorption coefficient, α α, from our two experimental measurements: τ 1 τ_1 (with the sample) and τ 2 τ_2 (without the sample).

The first bullet point says to combine our expressions for τ 1 τ_1 and τ 2 τ_2. The most convenient way to do this is to work with their inverses, the decay rates. Let's write them out:

1 τ 1 = c L (T + A + α L)

$$\frac{1}{\tau_1} = \frac{c}{L}(T + A + \alpha L)$$

1 τ 2 = c L (T + A)

$$\frac{1}{\tau_2} = \frac{c}{L}(T + A)$$

Now, watch what happens when we subtract the second equation from the first: 1 τ 1 − 1 τ 2 = $\frac{1}{\tau_1} - \frac{1}{\tau_2} =$ On the right side, the terms c L (T + A) $\frac{c}{L}(T + A)$ are common to both expressions, so they cancel out completely! We are left with just c L $\frac{c}{L}$ times α L αL. The L's cancel, and we are left with c times alpha.

So, our resulting equation is:

1 τ 1 − 1 τ 2 = c α

$$\frac{1}{\tau_1} - \frac{1}{\tau_2} = c\,\alpha$$

This is a beautiful and powerful result. We can now easily solve for alpha, which is the quantity we want to measure.

α = 1 c (1 τ 1 − 1 τ 2)

$$\alpha = \frac{1}{c}\left(\frac{1}{\tau_1} - \frac{1}{\tau_2}\right)$$

This is the central equation of Cavity Ring-Down Spectroscopy.

Now, let’s consider the profound “Advantages of this difference method,” highlighted in red.

Page 23:

Let's elaborate on those crucial advantages.

First, and most importantly, the mirror losses, T + A $T + A$, cancel out completely in the subtraction. Think about what this means. We don't need to know the absolute value of our mirror's reflectivity or its scattering losses. These can be very difficult to measure accurately. As long as these properties remain *stable* between our sample measurement (for τ 1 τ_1) and our reference measurement (for τ 2 τ_2), their exact value is irrelevant. This is a massive practical advantage, removing a huge source of systematic error.

Second, any intensity noise or fluctuation in the initial laser pulse, P 0 P_0, also cancels out. Why? Because P 0 P_0, or the related first pulse power P 1 P_1, only determines the initial amplitude of our decay curve. It appears as a pre-factor, not inside the exponential. When we perform a fit to the decay curve, we are fitting for the time constant *in the exponent*. The initial amplitude is just a scaling factor that doesn't affect the determination of τ τ. Since our final calculation for α α depends only on τ 1 τ_1 and τ 2 τ_2, the initial laser power has no bearing on the result. This is what makes CRDS so wonderfully immune to laser power noise.

These two points are the foundation of the technique's remarkable sensitivity and robustness.

Page 24

Now that we have our working equation for alpha, we need to understand how uncertainties in our measurements affect the uncertainty in our final result. This is a standard exercise in the propagation of measurement uncertainty.

Our goal, as stated, is to relate the uncertainties in our measured decay times, δ τ 1 $\delta\tau_1$ and δ τ 2 $\delta\tau_2$, to the resulting uncertainty in alpha, δ α $\delta\alpha$.

The standard formula for error propagation involves partial derivatives. We need to calculate the partial derivative of alpha with respect to tau_1, and the partial derivative of alpha with respect to tau_2. Our expression for alpha is α = 1 c (τ 1 − 1 − τ 2 − 1) $\alpha = \frac{1}{c}(\tau_1^{-1} - \tau_2^{-1})$. Let's take the derivatives.

The partial derivative of alpha with respect to tau_1, written d α d τ 1 $\frac{d\alpha}{d\tau_1}$, is 1 c × (− 1) × (τ 1 − 2) $\frac{1}{c} \times (-1) \times (\tau_1^{-2})$. So we get:

d α d τ 1 = − 1 c τ 1 2 .

$$\frac{d\alpha}{d\tau_1} = -\frac{1}{c\,\tau_1^2}.$$

Similarly, the partial derivative with respect to tau_2 is:

d α d τ 2 = 1 c τ 2 2 .

$$\frac{d\alpha}{d\tau_2} = \frac{1}{c\,\tau_2^2}.$$

Now, let's make a reasonable assumption for a practical experiment. We assume that our timing precision for both measurements is the same. That is, δ τ 1 $\delta\tau_1$ is approximately equal to δ τ 2 $\delta\tau_2$, and we can just call this value δ τ $\delta\tau$. We also know that for weak absorption, τ 1 τ_1 and τ 2 τ_2 will be very close to each other. So we can define an average ring-down time, τ τ, as:

τ = τ 1 + τ 2 2 .

$$\tau = \frac{\tau_1 + \tau_2}{2}.$$

This simplifies the analysis that follows.

Page 25:

Using these derivatives and the standard error propagation formula, which in its simplest form is (δ f) 2 = (d f d x) 2 (δ x) 2 + (d f d y) 2 (δ y) 2 ,

$$(\delta f)^2 = \left(\frac{df}{dx}\right)^2 (\delta x)^2 + \left(\frac{df}{dy}\right)^2 (\delta y)^2,$$

we can find the uncertainty in α α, which we call δ α $\delta\alpha$.

Under the condition of weak absorption, where | τ 1 − τ 2 | $|\tau_1 - \tau_2|$ is much, much smaller than the average time τ τ, the uncertainty formula simplifies significantly. The derivation is a bit tedious, but the result is shown on the slide: δ α ≈ 2 δ τ τ 3 | τ 1 − τ 2 | .

$$\delta\alpha \approx \frac{2\,\delta\tau}{\tau^3}\,|\tau_1 - \tau_2|.$$

Correction: A more standard propagation of error would yield (δ α) 2 = (1 c τ 1 2) 2 (δ τ 1) 2 + (1 c τ 2 2) 2 (δ τ 2) 2 .

$$(\delta\alpha)^2 = \left(\frac{1}{c\,\tau_1^2}\right)^2 (\delta\tau_1)^2 + \left(\frac{1}{c\,\tau_2^2}\right)^2 (\delta\tau_2)^2.$$

Assuming τ 1 τ_1 is close to τ 2 τ_2, this simplifies to δ α ≈ 2 δ τ c τ 2 .

$$\delta\alpha \approx \frac{\sqrt{2}\,\delta\tau}{c\,\tau^2}.$$

Let me re-evaluate the slide's formula. The presented formula δ α ≈ 2 δ τ τ 3 | τ 1 − τ 2 |

$$\delta\alpha \approx \frac{2\,\delta\tau}{\tau^3}\,|\tau_1 - \tau_2|$$

seems unusual. Let's re-derive. α = 1 c τ 2 − τ 1 τ 1 τ 2 ≈ 1 c τ 2 − τ 1 τ 2 .

$$\alpha = \frac{1}{c}\frac{\tau_2 - \tau_1}{\tau_1\,\tau_2} \approx \frac{1}{c}\frac{\tau_2 - \tau_1}{\tau^2}.$$

Then δ α = 1 c τ 2 | δ (τ 2 − τ 1) | .

$$\delta\alpha = \frac{1}{c\,\tau^2}\,|\delta(\tau_2 - \tau_1)|.$$

δ (τ 2 − τ 1) 2 = δ τ 2 2 + δ τ 1 2 = 2 δ τ 2 .

$$\delta(\tau_2 - \tau_1)^2 = \delta\tau_2^2 + \delta\tau_1^2 = 2\,\delta\tau^2.$$

So | δ (τ 2 − τ 1) | = 2 δ τ .

$$|\delta(\tau_2 - \tau_1)| = \sqrt{2}\,\delta\tau.$$

This leads to δ α ≈ 2 δ τ c τ 2 .

$$\delta\alpha \approx \frac{\sqrt{2}\,\delta\tau}{c\,\tau^2}.$$

The formula on the slide might be a different form or have a typo. However, I must interpret the slide as given. Let me proceed by explaining the implications of the formula on the slide, while being aware it might be non-standard.

Let's focus on the key takeaways from the uncertainty analysis, regardless of the exact form of the equation. The second bullet point is the most important message on this page. How can we improve our sensitivity? Improving sensitivity means making δ α $\delta\alpha$ as small as possible. Looking at the general dependence, we can see two clear pathways.

First, sensitivity improves with longer cavity ring-down times, τ τ. Since τ τ appears in the denominator (as τ 2 τ^2 or τ 3 τ^3), making τ τ larger will make δ α $\delta\alpha$ smaller. And how do we get a larger τ τ? By using mirrors with higher reflectivity, R R. This is the single most important factor in designing a sensitive CRDS system.

Second, sensitivity improves with better timing precision, δ τ $\delta\tau$. δ τ $\delta\tau$ represents the noise or uncertainty in our measurement of the decay time. It appears in the numerator, so making it smaller directly reduces our final uncertainty, δ α $\delta\alpha$. We achieve better timing precision by improving the signal-to-noise ratio (S / N S/N) of our decay curve measurement.

So, the recipe for high sensitivity is: Use the best mirrors you can get, and use a high-quality, low-noise detection system.

Page 26:

This slide translates our uncertainty analysis into concrete experimental design implications. We want to minimize $\Delta \alpha$ $\Delta\alpha$, so what should we do in the lab?

The first point reiterates our main conclusion: Increase the ring-down time, τ τ, by choosing mirrors with higher reflectivity, R R. This is the most direct way to gain orders of magnitude in sensitivity. However, there is a point of diminishing returns. As the slide notes, when you start using mirrors with reflectivities approaching unity (say, $R > 0.99995$ $R > 0.99995$), the mirror transmission, T T, becomes incredibly small. At this point, other intrinsic loss mechanisms, like diffraction and scatter (the A A term), can become comparable to or even larger than the transmission. When these losses start to dominate, simply increasing R R further doesn't give you the same dramatic improvement in ring-down time. You become limited by the quality of the mirror substrate and coating, not just its reflectivity.

The second major strategy is to enhance the signal-to-noise ratio of the measurement, which directly reduces the timing uncertainty, $\Delta \tau$ $\Delta\tau$. The slide lists several ways to do this:

1. Use low-noise detectors. A photodetector with lower intrinsic electronic noise will give a cleaner signal. 2. Average many decay traces. Since much of the noise is random, averaging hundreds or thousands of decay curves will cause the noise to average down, typically as the square root of the number of averages, while the exponential signal remains. This is a very powerful and common technique.

Page 27:

3. Electronic filtering. We can use appropriate electronic filters, either analog or digital, to remove noise components that are at frequencies far away from our signal's characteristic frequency content.

4. Suppressing mechanical vibrations. This is a crucial practical point. The ring-down time depends on the cavity length, L L. If the table is vibrating or there are acoustic noises in the room, the distance between the two mirrors can fluctuate slightly. This is called "cavity length jitter." This jitter can introduce noise into the decay traces, effectively worsening your timing precision. Therefore, building a mechanically stable and isolated cavity is essential for high-performance systems.

Finally, we have to consider an important experimental trade-off. A long ring-down time, τ τ, which is what we want for high sensitivity, necessarily means a slow decay. This means the data acquisition time for a single decay event increases. If τ τ is 100 microseconds, you need to record for several hundred microseconds to capture the full decay. This seems like it might slow down your experiment. However, modern pulsed lasers can often be fired at repetition rates of several kilohertz. This means that even if each decay takes a fraction of a millisecond, you can still acquire and average thousands of shots in just a few seconds. So, it's possible to have both a long ring-down time and rapid signal averaging.

Page 28:

This graph provides a powerful visual confirmation of our design principles. Let's break it down.

The plot shows the calculated Absorption Uncertainty, δ α $\delta\alpha$, on the vertical axis, versus the Mirror Reflectivity, R R, on the horizontal axis. Note that the vertical axis for δ α $\delta\alpha$ is a logarithmic scale, spanning from 10 − 10 10^{-10} up to 10 − 7 c m − 1 $10^{-7}\,\text{cm}^{-1}$. The horizontal axis for R R is linear, but covers only the very high reflectivity range from 0.990 0.990 to 1.000 1.000.

The graph was generated for a fixed set of parameters, noted in the top right: a cavity length L L of 50 c m $50\,\text{cm}$, and a timing precision δ τ $\delta\tau$ of 1 n s $1\,\text{ns}$.

The blue curve shows the resulting uncertainty, δ α $\delta\alpha$. The relationship is stark and dramatic. As we move from left to right—that is, as we increase the mirror reflectivity—the uncertainty in our measurement plummets by orders of magnitude.

At R = 0.990 $R = 0.990$, the uncertainty is around 10 − 7 10^{-7}. By the time we get to R = 0.999 $R = 0.999$, the uncertainty has dropped by nearly two orders of magnitude to about 2 × 10 − 9 2×10^{-9}. As we push R R even closer to 1, the curve becomes even steeper.

The annotation "delta alpha is proportional to 1 − R $1 - R$" captures the essence of this relationship. This comes directly from our uncertainty analysis: δ α $\delta\alpha$ is roughly proportional to 1 / τ 2 $1/\tau^2$, and τ τ itself is proportional to 1 / (1 − R) $1/(1 - R)$.

So, δ α $\delta\alpha$ goes roughly as (1 − R) 2 $(1 - R)^2$, which is an even stronger dependence, but the key takeaway is correct: as (1 − R) $(1 - R)$ gets smaller, so does the uncertainty. This plot unequivocally shows that the choice of mirrors is paramount for achieving high sensitivity.

Page 29:

Alright, let's solidify these concepts with a worked example. We'll start with a case of what we'll call "moderate" reflectivity.

We are given a set of parameters for our experiment.

First, the mirror reflectivity, R, is 0.999 0.999. This is often called "three nines" reflectivity. This means the total loss per reflection, 1 − R $1 - R$, is 10 − 3 10^{-3}.

The physical cavity length, L, is 1 m $1\,\mathrm{m}$. To be consistent with our other units, we should convert this to 100 cm $100\,\mathrm{cm}$.

The absorption coefficient, α α, of the trace gas we are trying to measure is 1 × 10 − 6 cm − 1 $1 \times 10^{-6}\,\mathrm{cm}^{-1}$. This is a very small absorption.

Finally, the speed of light, c, is 3.00 × 10 10 cm/s 3.00×10^{10} cm/s.

Our task is to compute the decay times with and without the sample, and then determine the uncertainty of our measurement.

Page 30:

Let's compute the decay times using the parameters from the previous slide.

First, let's calculate the decay time of the empty cavity, τ_2. We'll use the approximate formula τ 2 = L / c 1 − R

$$\tau_2 = \frac{L/c}{1 - R}$$

The term L / c L/c is 100 centimeters divided by (3 × 10 10 cm/s) $(3 \times 10^{10}\,\mathrm{cm/s})$, which is 3.33 nanoseconds. The denominator, 1 − R $1 - R$, is 10 − 3 10^{-3}. So, τ 2 = 3.33 × 10 − 9 seconds 10 − 3

$$\tau_2 = \frac{3.33 \times 10^{-9}\,\text{seconds}}{10^{-3}}$$

which equals 3.33 × 10 − 6 3.33×10^{-6} seconds, or three point three three microseconds.

Next, we compute the decay time with the absorber present, τ_1. The formula is τ 1 = L / c 1 − R + α L

$$\tau_1 = \frac{L/c}{1 - R + \alpha L}$$

The new term in the denominator is α L αL, which is (10 − 6 cm − 1) × (100 cm) = 10 − 4 $(10^{-6}\,\text{cm}^{-1}) \times (100\,\text{cm}) = 10^{-4}$. So the denominator is now 10 − 3 + 10 − 4 $10^{-3} + 10^{-4}$, which is 1.1 × 10 − 3 1.1×10^{-3}.

Let's re-calculate using the numbers on the slide. The denominator is 10 − 3 + 10 − 4 $10^{-3} + 10^{-4}$. The slide calculates τ_1 to be approximately three point zero three microseconds. Let's check: 100 cm / c 10 − 3 + 10 − 4 = 3.33 ns 1.1 × 10 − 3 = 3.03 microseconds .

$$\frac{100\,\text{cm}/c}{10^{-3} + 10^{-4}} = \frac{3.33\,\text{ns}}{1.1 \times 10^{-3}} = 3.03\,\text{microseconds.}$$

Yes, that's correct.

The difference in decay times, $\Delta\tau$, is τ_2 minus τ_1, which is 3.33 minus 3.03, giving zero point three zero microseconds. This is the small change in decay time that we have to measure.

Now for the uncertainty. Let's assume our experimental timing precision, $\delta\tau$, is plus or minus zero point zero three microseconds. This is ten percent of the difference we're trying to measure.

Using a simplified uncertainty formula, δ α ∝ δ τ c τ 2 .

$$\delta\alpha \propto \frac{\delta\tau}{c\,\tau^2}.$$

Plugging in the numbers, the slide calculates the absolute uncertainty, δα, to be approximately 4 × 10 − 7 4×10^{-7} inverse centimeters.

How does this compare to the true value of α we were trying to measure? The relative uncertainty is δα divided by α. That's (4 × 10 − 7) / (1 × 10 − 6) $(4 \times 10^{-7})/(1 \times 10^{-6})$, which is 0.4, or a whopping 40 percent!

Page 31:

The conclusion from this first worked example is clear and sobering.

Using mirrors with a reflectivity of 0.999 0.999 is insufficient for making a high-accuracy measurement of this particular trace species. A measurement with a 40 % 40% uncertainty is not very useful for quantitative analysis. The change in the decay time caused by the absorber was simply too small compared to our ability to measure that change. To do better, we need to find a way to make the difference between τ 1 τ_1 and τ 2 τ_2 much larger. And as we know, the way to do that is to use better mirrors.

Page 32:

Let's now repeat the exact same worked example, but this time we'll upgrade our hardware. We'll perform the measurement with high-reflectivity mirrors.

Our new reflectivity, R R, is zero point nine nine nine nine, or "four nines." This means the intrinsic loss per reflection, 1 − R $1 - R$, is now ten times smaller, at

10 − 4 10^{-4}. All other parameters (L L, α α, c c, and our timing precision δ τ $\delta\tau$) remain the same.

Let's re-calculate the decay times.

The empty cavity decay time, τ 2 τ_2, is (L / c) / (1 − R) $(L/c)/(1-R)$, which is (100 c m / c) / (10 − 4) $(100\ \mathrm{cm}/c)/(10^{-4})$. This is approximately 33 μ s $33\ \mu s$. Notice that just by adding one 'nine' to the reflectivity, our ring-down time has increased by a factor of ten!

Now for the decay time with the absorber, τ 1 τ_1. The denominator is (1 − R + α L) $(1-R+\alpha L)$, which is 10 − 4 + 10 − 4 = 2 × 10 − 4 $10^{-4}+10^{-4}=2\times 10^{-4}$. So τ 1 τ_1 is (100 c m / c) / (2 × 10 − 4) $(100\ \mathrm{cm}/c)/(2\times 10^{-4})$, which is approximately 16.5 μ s $16.5\ \mu s$.

The difference in decay times, Δ τ $\Delta\tau$, is now 33 − 16.5 $33-16.5$, which is 16.5 μ s $16.5\ \mu s$. This is a huge improvement! The difference we need to measure is now 16.5 μ s $16.5\ \mu s$, not the tiny 0.3 μ s $0.3\ \mu s$ from before.

We are still using the same detection system, so our timing precision, δ τ $\delta\tau$, remains ± 0.03 μ s $\pm 0.03\ \mu s$.

Page 33:

Now let’s calculate the final uncertainty for the high-reflectivity case.

With a much larger difference between τ 1 τ_1 and τ 2 τ_2, and much longer decay times overall, our uncertainty formula predicts a much smaller δ α $\delta\alpha$. Plugging in the new values, the calculation yields:

delta alpha is approximately 6 × 10 − 9 c m − 1 $6\times 10^{-9}\ \mathrm{cm}^{-1}$.

Let's compare this to the true value of α, which was 1×10^{-6}.

The relative uncertainty, $\delta\alpha/\alpha$, is now $(6 \times 10^{-9})/(1 \times 10^{-6})$, which is 0.006, or just zero point six percent.

The conclusion is powerful and unambiguous. By moving from 99.9% to 99.99% reflectivity, we have improved our measurement accuracy from a useless 40% uncertainty to a very respectable 0.6% uncertainty. A high-finesse cavity dramatically enhances detection accuracy. This pair of examples perfectly illustrates the core principle of CRDS.

Page 34:

These two plots provide a perfect visual summary of the two worked examples we just went through. They show exactly why high-reflectivity mirrors are so crucial.

On both plots, the vertical axis is the normalized intensity on a logarithmic scale, and the horizontal axis is time in microseconds. On a log-linear plot, an exponential decay becomes a straight line, and the slope of that line is related to the decay time. A steeper slope means a shorter decay time.

Let's look at the left panel, for the "Moderate Reflectivity" case, $R = 0.999$. The blue line represents the decay of the empty cavity, with its time constant τ_2 of $3.3\,\mu s$. The orange line shows the decay with the absorbing sample, with its shorter time constant τ_1 of $3.03\,\mu s$. As you can see, the two lines are very close together. Their slopes are almost identical. The difference between them, $\Delta\tau$, is only $0.30\,\mu s$. If your measurement has any noise, telling these two lines apart becomes very difficult.

Now look at the right panel, for the "High Reflectivity" case, R = 0.9999 $R = 0.9999$. The first thing you'll notice is that the time axis is much longer, going out to 60 μ s $60\,\mu s$. The decays are much, much slower. The blue line for the empty cavity has a time constant of 33 μ s $33\,\mu s$, and the orange line for the sample has a time constant of 16.5 μ s $16.5\,\mu s$. Now, the two lines are clearly and widely separated. The difference, Δ τ $\Delta\tau$, is a large 16.8 μ s $16.8\,\mu s$. This difference is vastly larger than our assumed timing uncertainty of 0.03 μ s $0.03\,\mu s$, making a precise and accurate measurement not just possible, but straightforward. This is a beautiful illustration of how finesse enhances sensitivity.

Page 35:

Let's now formally introduce two important figures of merit that are often used to characterize a CRDS system: the effective absorption path length and the cavity finesse.

First, as we've discussed, the incredible sensitivity of CRDS comes from the fact that the light makes many, many passes through the sample. We can calculate the *effective number of passes* the light makes through the sample, N_pass.

In the simple case of an empty cavity, this is approximately given by the formula:

N p a s s = 1 1 − R

$$N_{\text{pass}} = \frac{1}{1 - R}$$

So for our R = 0.9999 $R = 0.9999$ mirror, the light makes 1 10 − 4 = 10,000 $\frac{1}{10^{-4}} = 10{,}000$ passes on average before being lost.

From this, we can calculate the *effective absorption length*, L_eff. It's simply the physical length of the cavity, L, multiplied by the number of passes. However, since the light traverses the length L twice per round trip, a more careful derivation gives the formula shown:

L e f f = L 1 − R

$$L_{\text{eff}} = \frac{L}{1 - R}$$

This formula gives us the length of an equivalent single-pass absorption cell that would yield the same total absorption.

Page 36:

Let's plug in the numbers from our high-reflectivity example. If the physical cavity length L L is one meter, and the reflectivity R R is 0.9999, then the effective path length, L e f f L_{eff}, is 1 meter divided by (1 − 0.9999) $(1 - 0.9999)$, which is 1 meter divided by 10 − 4 10^{-4}. This gives an L e f f L_{eff} of 10,000 meters, or ten kilometers. This number should be burned into your memory; it perfectly encapsulates the power of this technique.

Next, we introduce a closely related and very important parameter: the Cavity Finesse, denoted by a script capital F. The finesse is a dimensionless quantity that characterizes the quality of an optical resonator. It's defined as:

F = π R 1 − R .

$$\mathcal{F} = \frac{\pi\sqrt{R}}{1 - R}.$$

Since R R is very close to 1, the square root of R R is also very close to 1, so a common approximation is

F ≈ π 1 − R .

$$\mathcal{F} \approx \frac{\pi}{1 - R}.$$

You can see that Finesse, just like the ring-down time and the effective path length, is inversely proportional to the cavity losses (1 − R) $(1 - R)$. A high-finesse cavity is a low-loss cavity. In fact, finesse is directly proportional to the ring-down time and therefore to the sensitivity of the measurement. A finesse of over 30,000 is common for CRDS.

The final point is an important practical one. Our simple formulas assume losses are only from transmission and absorption. But diffraction can also be a loss. To achieve the highest possible finesse, one must ensure that the input laser beam is properly "mode-matched" to the cavity. This means using a system of lenses to focus the laser beam so that its waist size and curvature perfectly match the cavity's fundamental TEM-zero-zero eigenmode. This excites the lowest-loss mode of the cavity and minimizes any additional losses due to diffraction, which would otherwise limit the achievable finesse.

Page 37

Page 38

Let's now move on to discuss a typical experimental setup for performing Pulsed CRDS. What are the key hardware components you would find in a real-world system?

First, you obviously need a laser. For pulsed CRDS, you need a short-pulse, tunable laser. A very common choice, especially in research labs, is a dye laser or an Optical Parametric Oscillator (OPO), which is pumped by a high-power pulsed laser like a Q-switched Nd:YAG laser. The tunability is essential for scanning across a molecular absorption feature to record a spectrum.

Next, you need mode-matching optics. As we just discussed, this typically consists of a combination of lenses and apertures. Their job is to take the output beam from the laser and shape it perfectly to align its beam waist with the cavity's fundamental eigenmode, ensuring efficient coupling of light into the cavity and minimizing diffraction losses.

Of course, the heart of the system is the cavity itself, formed by two high-reflectivity mirrors. In many practical setups, these mirrors also serve as the windows for the absorption cell or vacuum chamber that contains the gas sample.

Finally, you need the detection system. This consists of a fast photodetector (like a photomultiplier tube or an avalanche photodiode) to convert the weak, fast leakage pulses into an electrical signal. This signal is then digitized. A traditional instrument for this is a "boxcar integrator," which is specifically designed to sample a signal in a very narrow time window and is good for averaging. More commonly today, one would use a fast digital oscilloscope or a dedicated high-speed digitizer card.

Page 39:

There are a couple of important temporal requirements to consider in a pulsed CRDS experiment.

First, there's a condition on the laser pulse duration, which we'll call T p T_p. For the cleanest measurement, where you can actually resolve the individual leakage pulses coming out of the cavity, the pulse duration T p T_p should be shorter than the cavity round-trip time, T R T_R. Remember, T R T_R is typically a few nanoseconds. So this requires using a nanosecond or picosecond laser. Seeing clearly separated leakage pulses can sometimes help in making a more accurate exponential fit, as you have distinct data points along the decay curve.

However, this is not a strict requirement. What happens if your laser pulse is longer than the round-trip time? For example, if T p T_p is greater than T R T_R. In this case, before the first part of the pulse has even completed a round trip, more of the pulse is still entering the cavity. The result is that the individual leakage pulses overlap in time. But this is perfectly fine! You will no longer see a train of discrete pulses, but you will still observe a smooth, continuous decay envelope after the laser pulse turns off. Crucially, this decay envelope is still a pure exponential, and its time constant, τ τ, is still determined by the total cavity losses in exactly the same way. So the technique works perfectly well in this regime too.

Page 40:

This diagram provides an excellent schematic of a complete Pulsed Cavity Ring-Down Spectroscopy setup. Let's trace the signal path from start to finish.

On the far left, we have the "Pulsed Laser," our source of light.

The beam emerges and passes through a set of mode-matching optics, here represented by two lenses, L1 and L2, and an aperture between them. This system, as we know, shapes the beam for optimal injection into the cavity.

The beam then enters the "Sample Cell," which is the region between the two high-reflectivity mirrors, M1 and M2. The mirrors are specified as having reflectivity of approximately 99.99%. Inside this cell, the light pulse is "trapped," bouncing back and forth as shown by the multiple red lines. A gas inlet and outlet, which are not shown, would allow for the introduction and removal of the sample gas.

With each bounce on mirror M2, a small amount of light leaks out. This "Leakage" light is directed onto a "Photodetector."

The photodetector converts the decaying light signal into a voltage signal. This electrical signal is then sent to the "Data Acquisition / Oscilloscope" system.

An inset shows what the recorded signal looks like. The vertical axis is Intensity and the horizontal axis is Time. We see a classic exponential decay curve. The data acquisition system's job is to digitize this curve and then use a computer to fit it to the mathematical function

I (t) = I 0 e − t / τ

$$I(t) = I_0\, e^{-t/\tau}$$

The result of this fit is the all-important number: the ring-down time, τ τ.

Page 41:

While pulsed CRDS is very common, there's an important and powerful variant of the technique that uses a continuous-wave, or CW, laser. This is known as CW-CRDS.

In this scheme, you don't use a short pulse. Instead, you illuminate the cavity with a highly stabilized, narrow-linewidth continuous-wave laser. You tune the laser frequency (or the cavity length) so that it is exactly on resonance with one of the cavity modes. This causes a large amount of optical power to build up inside the cavity.

Then, at a time we'll call t = 0 $t = 0$, the input laser beam is abruptly switched off. This is typically done with a very fast optical switch, like an Acousto-Optic Modulator (AOM) or an Electro-Optic Modulator (EOM), which can turn the beam off in tens of nanoseconds.

Once the input is cut off, the energy that was stored inside the cavity has nowhere to go but to leak out through the mirrors. And just like in the pulsed case, this stored intra-cavity field decays exponentially with the exact same time constants, τ 1 τ_1 or τ 2 τ_2, that we derived before.

This CW approach has some significant advantages over pulsed excitation. First, CW lasers generally have much narrower spectral linewidths than pulsed lasers. This allows for much higher spectral resolution, which is critical if you want to resolve very fine details in an absorption spectrum. Second, you don't need high-peak-power pulsed lasers. High peak powers can sometimes lead to undesirable effects like sample saturation, or even damage to the optics. CW-CRDS avoids this.

Page 42:

Continuing with the advantages of the continuous-wave variant.

CW sources are generally more stable and have lower amplitude noise than many high-power pulsed laser systems. While standard CRDS is theoretically immune to amplitude noise, in practice, extreme shot-to-shot fluctuations in pulsed lasers can still present challenges for the detection electronics. The stability of CW lasers can lead to cleaner decay signals.

A final, more advanced point is that the use of a stable, single-frequency CW laser is the key to enabling some of the most sensitive detection methods, such as the heterodyne detection schemes we will touch on later. These methods rely on mixing the signal with a stable local oscillator, which is much more straightforward to implement with a CW laser.

So, in summary, while the basic principle of measuring a decay time remains the same, the choice between pulsed and CW excitation depends on the specific requirements of the experiment, such as the desired spectral resolution, the robustness of the sample, and the ultimate sensitivity required.

Page 43:

Let's return to a practical issue we touched upon earlier: the phenomenon of interference or "beat" phenomena, and the concept of mode averaging.

This problem arises in pulsed CRDS when the laser pulse is broad enough in frequency to excite multiple cavity modes simultaneously. Let's say it excites two adjacent longitudinal modes. These two modes have slightly different frequencies. When they both leak out of the cavity and onto the detector, they interfere with each other. This interference creates a sinusoidal "beat note" in the detected intensity, whose frequency is equal to the difference in frequency between the two modes.

This beat note is then superimposed on top of the overall exponential decay. The result is a signal that looks like a decaying sine wave, not a pure exponential. This modulation can corrupt the fitting process and lead to an inaccurate determination of the decay time, τ τ.

So how do we deal with this? The second bullet point provides the most common solution. The exact phase of the beat note depends on the relative phase of the different laser frequency components that excited the modes. For most pulsed lasers, this relative phase is random from one laser shot to the next. This means that on one shot the beat note might be a sine wave, on the next a cosine wave, and on the next something in between.

If we average many of these decay traces together, these random-phase beat notes will average out to zero. The underlying pure exponential decay, however, is the same for every shot and adds up coherently. The result is that averaging reveals the pure exponential, free from the mode-beating artifacts.

Page 44:

While averaging is an effective way to remove mode beats, there is a more elegant solution, particularly for experiments that require the absolute highest spectral resolution.

This is single-mode excitation. If you can ensure that your laser only excites one single cavity mode, then there are no other modes for it to interfere with, and the problem of beat notes is avoided entirely.

How is this achieved? It requires two things. First, you need a laser with a very narrow linewidth, one that is significantly narrower than the spacing between the cavity modes (the Free Spectral Range). Second, you need to actively tune the cavity length (for example, with a piezo actuator on one mirror) or the laser frequency to ensure that the single laser frequency stays precisely locked to the peak of the single cavity resonance. This is the principle behind many high-performance CW-CRDS systems.

This approach is especially important for high-resolution spectroscopy. Imagine you are trying to measure the precise shape of a very narrow absorption feature. If your signal is contaminated by mode beats, the beat structure could be mistaken for, or completely mask, the fine absorption features you are trying to see. Single-mode excitation provides a clean, artifact-free decay, allowing the true spectral lineshape to be recovered.

Page 45:

So, we now have a robust method for measuring the absorption coefficient, α α, at a single, specific laser wavelength. But the real power of spectroscopy is in measuring how α α changes as a function of wavelength to reveal an absorption *spectrum*. How do we do this with CRDS?

The process is conceptually simple, as outlined here. First, we use our tunable laser to scan its wavelength, λ λ, across the molecular absorption region of interest. For example, we might scan across a vibrational overtone band of a molecule.

At each discrete wavelength step in our scan, we perform a complete CRDS measurement. This means we measure two things: the ring-down time with our sample in the cavity, which gives us τ 1 (λ) $\tau_1(\lambda)$, and a reference ring-down time, τ 2 τ_2. The reference measurement can be done by measuring the decay time in a truly empty cavity, or, more conveniently, by taking a measurement at an "off-line" wavelength—a wavelength nearby where we know the molecule does not absorb. This τ 2 τ_2 measurement characterizes the baseline losses of the cavity itself.

Then, at each wavelength, we compute the difference using our main formula. Let's define a quantity Δ $\mathit{\Delta}$, which is a function of λ λ, as $1 / \tau 1 (\lambda) - 1 / \tau 2$ $1/\tau_1(\lambda) - 1/\tau_2$. As we know from our derivation, this quantity, Δ $\mathit{\Delta}$, is approximately equal to $c \alpha (\lambda)$ $c\,\alpha(\lambda)$. So by calculating this difference, we directly determine the absorption coefficient at that specific wavelength.

Page 46:

Once we have calculated $\alpha (\lambda)$ $\alpha(\lambda)$ at every point in our wavelength scan, the final step is straightforward. We simply plot our calculated absorption coefficient, α α, as a function of wavelength or, more commonly in spectroscopy, as a function of wavenumber. The result is a high-resolution absorption spectrum of our sample.

Depending on the conditions in our sample cell (primarily the pressure and temperature), the spectral lines we observe will have a certain shape. At low pressures, the lineshape is typically determined by the thermal motion of the molecules, leading to a "Doppler-limited" Gaussian profile. At higher pressures, collisions between molecules become significant, leading to "pressure-broadened" Lorentzian lineshapes.

The slide gives a concrete example: measuring an overtone band of the molecule HCN, or hydrogen cyanide. These overtone transitions are typically very weak, making them a perfect application for CRDS. The resulting spectrum would show the characteristic pattern of rotational lines within the vibrational band.

Page 47:

And here is a beautiful example of exactly such a spectrum. This plot shows an experimental CRDS absorption spectrum of an overtone band of HCN.

Let's examine the axes. The horizontal axis is wavenumber in units of inverse centimeters, spanning from about 6467 to 6577. The vertical axis is the absorption coefficient, α α, in arbitrary units, normalized to a maximum of one.

What we see is a classic ro-vibrational spectrum of a linear molecule. It consists of two main "branches" of lines. On the left, at lower wavenumbers, we have the P-branch, where the rotational quantum number decreases during the vibrational transition. You can see the individual lines labeled P(2), P(4), P(6), and so on. The spacing between these lines is characteristic of the molecule's rotational constant.

On the right, at higher wavenumbers, we have the R-branch, where the rotational quantum number increases. These lines are labeled R(0), R(2), R(4), etc. There is a gap in the center where the Q-branch (with no change in rotational quantum number) would be, but it is forbidden for this type of transition in HCN.

The ability to resolve each of these individual rotational lines with such a high signal-to-noise ratio, despite the transition being intrinsically weak, is a testament to the power and resolution of Cavity Ring-Down Spectroscopy.

Page 48:

Let's take a moment to summarize the key advantages of CRDS that we've discussed so far. This really brings home why the technique is so powerful.

First, and perhaps most importantly, is its insensitivity to laser intensity noise. I can't stress this enough. Because we are fitting a decay *constant*, not a signal *amplitude*, fluctuations in the laser shot-to-shot energy largely do not affect the result. This overcomes the primary limitation of traditional absorption spectroscopy.

Second, the use of high repetition rate lasers, typically operating in the kilohertz range, allows for extremely rapid signal averaging. Noise in the measurement tends to be random, so by averaging thousands of decay traces, we can improve the signal-to-noise ratio dramatically—typically by the square root of the number of averages—in just a matter of seconds.

Third, the technique provides an extremely long effective path length. As we saw in our example, a simple one-meter cavity can provide an effective interaction length of up to 10 4 10^4 times that of a single pass, which is ten kilometers. This massive amplification is what allows us to see incredibly weak absorption signals.

Page 49:

Continuing our summary of advantages:

The fourth key advantage is the simplicity of the data analysis. In the ideal case, the signal is a pure exponential decay. Fitting an exponential function to data is a mathematically robust and straightforward procedure. There is a single parameter, τ τ, to extract. This means that to determine the concentration of a species, which is proportional to the absorption coefficient α α, we don't need to perform any complex lineshape deconvolution or have a perfect theoretical model of the line profile. The determination of the total absorption is direct and simple.

These four points—immunity to intensity noise, rapid averaging, long effective path length, and simple analysis—combine to make CRDS one of the most sensitive, robust, and versatile techniques available for quantitative absorption spectroscopy.

Page 50:

Of course, no technique is without its limitations and practical conditions that must be met for optimal performance. Let's discuss some of these for CRDS.

The most significant limitation is a technological one: CRDS requires mirrors with very, very high reflectivity. As we saw in our worked examples, to truly outperform other sensitive methods, you typically need mirrors with reflectivities of R ≥ 0.9995 $R \geq 0.9995$, and often much higher. Producing such high-quality mirrors can be challenging and expensive, and they may not be readily available for all wavelength regions, particularly in the ultraviolet or far-infrared.

Next, there are some conditions relating the cavity properties to the laser and absorber properties that are important for achieving high spectral resolution. The two inequalities shown here relate the various bandwidths, expressed in terms of angular frequency, ω ω.

The first inequality is: δ ω R $\delta\omega_{\mathrm{R}}$, which is the linewidth of the cavity resonance itself, should be less than δ ω a $\delta\omega_{\mathrm{a}}$, the linewidth of the absorber. The cavity resonance linewidth is given by one over the ring-down time, or in this notation, 1 T R $\frac{1}{T_{\mathrm{R}}}$ is actually the Free Spectral Range. The cavity mode linewidth is F S R Finesse $\frac{\mathrm{FSR}}{\mathrm{Finesse}}$. Let's assume δ ω R $\delta\omega_{\mathrm{R}}$ refers to the cavity mode linewidth. This condition means the instrument's resolution should be finer than the feature you're trying to resolve.

The second inequality is: δ ω L $\delta\omega_{\mathrm{L}}$, the laser linewidth, should also be less than δ ω a $\delta\omega_{\mathrm{a}}$. Again, your light source must be sharper than the feature you wish to measure.

These conditions ensure that the measured spectrum is a true representation of the absorber's lineshape, not one that is artificially broadened by the instrument itself.

Page 51:

Here are a few more important practical considerations.

The first point is a subtle but important one regarding potential spectroscopic complications. The cavity relaxation time, which is our ring-down time τ τ, must exceed the excited-state lifetime of the absorbing molecule, which is denoted T e x c T_{exc}. Why is this? If the molecule stays in the excited state for a long time, and the light field in the cavity is intense, an excited molecule can be hit by another photon and be stimulated to emit its energy back into the light field. This stimulated emission effectively "cancels out" an absorption event. This can lead to a non-linear relationship between the measured loss and the true concentration, a phenomenon known as saturation. By ensuring the cavity field decays quickly compared to the molecular lifetime, we can largely avoid these complications and stay in the linear absorption regime described by the Beer-Lambert law.

Second, as we've discussed, the discrete mode structure of the cavity can be a problem. The resonance peaks of the cavity can overlay the absorption spectrum we are trying to measure, potentially distorting it. There are two main ways to mitigate this. For broadband detection using a pulsed laser, we can average over many laser shots to smooth out the mode structure. For high-resolution CW experiments, the solution is to synchronously tune the cavity length as the laser frequency is scanned. This keeps the single cavity mode being used perfectly on resonance with the laser, effectively making the cavity "invisible" and allowing the true absorption spectrum to be traced out.

Page 52:

Let's consider another numerical example, this time focusing on the interplay between the various bandwidths involved in a measurement. These are the mode-structure considerations.

We are given the following parameters for a pulsed CRDS experiment: The cavity length, L L, is 0.5 m $0.5\,\mathrm{m}$, or 50 c m $50\,\mathrm{cm}$. From this, we can immediately calculate the round-trip time, T R T_{R}, which is 2 L / c $2L/c$. This comes out to 3.3 × 10 − 7 s 3.3×10^{-7} s.

The mirror reflectivity, R R, is 0.995 0.995. This is a moderately good mirror, but not exceptional.

From these values, we can calculate the cavity bandwidth, or more precisely, the linewidth of an individual cavity resonance, δ ω R $\delta\omega_{\mathrm{R}}$. This is given by the Free Spectral Range divided by the Finesse. A simpler, related quantity is 1 / T R $1/T_{\mathrm{R}}$, which is the Free Spectral Range. A more accurate representation of the cavity linewidth is 1 / τ $1/\tau$. Let's assume the slide meant the cavity mode linewidth is on the order of 1 / τ $1/\tau$.

The empty cavity decay time, τ τ, would be τ = L / c 1 − R

$$\tau = \frac{L/c}{1-R}$$

which is τ = 1.67 n s 0.005 = 3.3 × 10 − 7 s .

$$\tau = \frac{1.67\ \mathrm{ns}}{0.005} = 3.3 \times 10^{-7}\ \mathrm{s}.$$

So 1 / τ $1/\tau$ is about 3 × 10 6 3×10^{6} radians per second, or 3 × 10 6 3×10^{6} per second, as written. So, δ ω R $\delta\omega_{\mathrm{R}}$ here represents the cavity resonance linewidth.

Page 53:

Now let's bring in the other players: the laser and the molecule.

We are told the laser pulse has a duration of 10 − 8 10^{-8} seconds, or 10 nanoseconds. Due to the Fourier transform relationship, a pulse of this duration has a minimum possible bandwidth, known as the Fourier-limit. This laser bandwidth, δ ω L $\delta\omega_{\mathrm{L}}$, is on the order of the inverse of the pulse duration, which is 10 8 s − 1 $10^{8}\ \mathrm{s}^{-1}$.

Let's compare this to the cavity linewidth. The laser bandwidth, 10 8 s − 1 $10^{8}\ \mathrm{s}^{-1}$, is much greater than the cavity mode linewidth, 3 × 10 6 s − 1 $3 \times 10^{6}\ \mathrm{s}^{-1}$. This confirms that this broad laser pulse will indeed excite many cavity modes simultaneously.

Wait, the slide says "single longitudinal mode is excited". This must be a typo, it should say "multiple longitudinal modes are excited". Let me re-read. Ah, perhaps the logic is different. "Because δ ω L > δ ω R $\delta\omega_{\mathrm{L}} > \delta\omega_{\mathrm{R}}$, single longitudinal mode is excited". This seems backwards. Let's reconsider. Maybe δ ω R $\delta\omega_{\mathrm{R}}$ is the FSR.

F S R = c 2 L = 3 × 10 8 2 ⋅ 0.5 = 3 × 10 8 r a d / s .

$$\mathrm{FSR} = \frac{c}{2L} = \frac{3 \times 10^{8}}{2 \cdot 0.5} = 3 \times 10^{8}\ \mathrm{rad/s}.$$

Laser bandwidth is 10 8 s − 1 $10^{8}\ \mathrm{s}^{-1}$. So laser bandwidth < FSR. In this case, it is possible to excite only a single mode if the laser frequency is tuned precisely to it. Let's assume this interpretation. So, the laser is narrow enough to excite only one mode at a time.

But now we introduce the absorber. A typical Doppler-broadened absorption width for a molecule in the visible spectrum is around 1.5 × 10 9 1.5×10^{9} radians per second, which is about 240 MHz. Let me re-check my numbers. A typical Doppler width is about 1 GHz, which is 2 π × 10 9 $2\pi \times 10^{9}$ rad/s. The number on the slide is 1.5 × 10 7 s − 1 $1.5 \times 10^{7}\ \mathrm{s}^{-1}$, which is quite small, about 2.4 MHz. Let me assume the numbers on the slide are correct for the context. This Doppler width, 1.5 × 10 7 s − 1 $1.5 \times 10^{7}\ \mathrm{s}^{-1}$, exceeds both the laser bandwidth (10^8? No, the laser is wider here) and the cavity linewidth. Let me restart the interpretation of this slide, as it seems inconsistent.

Alternative interpretation: Let's assume there is a typo in the inequality sign or the conclusion. Given: Cavity linewidth δ ω R ≈ 3 × 10 6 s − 1 $\delta\omega_{\mathrm{R}} \approx 3 \times 10^{6}\ \mathrm{s}^{-1}$. Laser bandwidth δ ω L ≈ 10 8 s − 1 $\delta\omega_{\mathrm{L}} \approx 10^{8}\ \mathrm{s}^{-1}$. Clearly, δ ω L > δ ω R $\delta\omega_{\mathrm{L}} > \delta\omega_{\mathrm{R}}$. This means the laser is broad enough to cover and excite *many* cavity modes. Let's assume the slide's text "single longitudinal mode is excited" is an error.

Now, let's consider the Doppler width of the absorber, given as ≈ 1.5 × 10 9 s − 1 $\approx 1.5 \times 10^{9}\ \mathrm{s}^{-1}$ (a more typical 240 MHz value). This is much wider than both the laser bandwidth and the cavity mode linewidth. This means that the ultimate spectral resolution of the experiment will not be limited by our instrument (the laser or the cavity), but rather by the intrinsic properties of the sample itself (the Doppler broadening). This is a good thing — we are truly measuring the molecule's spectrum.

The final implication is critical. If we are scanning the laser wavelength to trace out this absorption feature, we must ensure that one of the cavity modes remains on resonance with the laser. Since the laser is narrow enough to fit within the Doppler profile, we need to tune the cavity length synchronously with the laser wavelength. This keeps a cavity mode "following" the laser as it scans, allowing for a continuous, high-resolution measurement of the absorption profile.

Page 54:

We've seen that CRDS is incredibly sensitive. But can we do even better? For the most demanding applications, the ultimate sensitivity of a single-shot measurement is often limited by two technical noise sources: residual electronic noise in the detector and its amplifier, and mechanical cavity-length jitter, which we discussed earlier.

So, is there a way to overcome these technical noise limits? Yes, there is. The solution is a sophisticated technique called optical heterodyne detection.

The core idea of heterodyne detection is to mix our weak signal of interest—in this case, the decaying light from the cavity—with a strong, stable reference beam, which is called the "local oscillator" or LO. The mixing occurs on the photodetector. As we'll see, the interference between the weak signal and the strong LO amplifies the signal and shifts it to a radio frequency, moving it away from the low-frequency noise that often plagues direct detection measurements.

The rest of the slide provides an outline of how such an experiment is typically implemented.

Page 55:

Let's walk through the experimental outline for a heterodyne-detected CRDS system.

First, you start with a single-mode, highly stable laser, like a diode laser. The beam from this laser is split into two paths using a beam splitter.

One path serves as the strong local oscillator, or LO. A key feature of this technique is that this LO beam is actively locked to one of the cavity modes, ensuring it is always on resonance and has a stable phase relationship with the cavity.

The second path becomes the "probe beam." This is the beam that will actually measure the absorption. This beam is first sent through an acousto-optic modulator (AOM). The AOM does two things. First, it frequency-shifts the light by an amount precisely equal to one free spectral range of the cavity. This makes the probe beam resonant with the cavity mode *adjacent* to the one the LO is locked to. Second, the AOM can be used to rapidly modulate the intensity of the probe beam, for example at a frequency of 40 kilohertz as suggested here. This intensity modulation is what we will ultimately detect.

Finally, the strong, stable LO beam and the weaker, intensity-modulated probe beam are recombined on a second beam splitter. These two co-propagating fields then traverse the cavity together and fall onto the detector. At the detector, they interfere, or "beat," creating the heterodyne signal that contains our information.

Page 56:

This diagram shows a detailed schematic of a Heterodyne-Detected Cavity Ring-Down Spectroscopy setup. It looks complex, but let's trace the signals.

We start at the top left with a "Single-mode Diode Laser." The beam is split at the first beam splitter, BS1.

The transmitted beam is our "Probe Beam." It goes through an "AOM," which is driven by an "RF Driver" that both shifts its frequency by one FSR and gets modulated by a 40 kHz signal from a "Modulator."

The reflected beam from BS1 is our "Local Oscillator" or LO. It is marked "Locked to Cavity Mode."

The probe and LO beams are recombined at the second beam splitter, BS2, and sent into the "High-Finesse Cavity" (M1, M2).

The light leaking out of the cavity goes through a "Polarizer" (P) and hits the "Photodiode." The photodiode detects the beat note between the probe and the LO.

Now look at the feedback loops. The signal from the photodiode is sent to a "PDH Lock" system. Pound-Drever-Hall locking is a standard technique to generate an error signal that is used to provide "Feedback to lock LO," by adjusting the laser frequency or cavity length, keeping the LO perfectly on resonance.

The photodiode signal also contains our 40 kHz heterodyne signal. This signal is sent to a "Mixer," which also receives a 40 kHz reference signal. The output of the mixer is sent to a "Lock-in Amp," which demodulates the signal, extracting the amplitude and phase of the heterodyne beat note with incredible sensitivity. This is our final measurement signal.

Page 57:

Let's look at the mathematics behind the heterodyne signal to understand why it's so powerful.

The total electric field at the detector is the sum of the signal field, E s (t) $E_s(t)$, and the local oscillator field, E L O E_{LO}. The LO field has a slight frequency offset, δ ν $\delta\nu$, and a phase ϕ ϕ relative to the signal. The detector measures intensity, which is proportional to the absolute square of the total electric field.

So, the total intensity, I T (t) $I_{\mathrm{T}}(t)$, is proportional to the magnitude squared of (E s (t) + E L O e i (2 π δ ν t + ϕ)) $\left(E_{\mathrm{s}}(t) + E_{LO}\, e^{i(2\pi\delta\nu t+\phi)}\right)$.

I T (t) ∝ | E s (t) + E L O e i (2 π δ ν t + ϕ) | 2 .

$$I_{\mathrm{T}}(t) \propto |E_{\mathrm{s}}(t) + E_{LO}\, e^{i(2\pi\delta\nu t+\phi)}|^2.$$

When we expand this squared term, we get three terms, as shown in the second equation: I_T(t) equals the magnitude of E s (t) $E_{\mathrm{s}}(t)$ squared, plus the magnitude of E L O E_{LO} squared, plus a cross term: 2 times E s (t) $E_{\mathrm{s}}(t)$ times E L O E_{LO} times the cosine of (2 π δ ν t + ϕ) $(2\pi\delta\nu t + \phi)$.

I T (t) = | E s (t) | 2 + | E L O | 2 + 2 E s (t) E L O cos (2 π δ ν t + ϕ) .

$$I_{\mathrm{T}}(t) = |E_{\mathrm{s}}(t)|^2 \; + \; |E_{LO}|^2 \; + \; 2\, E_{\mathrm{s}}(t)\, E_{LO} \cos(2\pi\delta\nu t + \phi).$$

Let's analyze these terms. | E s | 2 $|E_{\mathrm{s}}|^2$ is the signal we would measure in direct detection; it's very weak. | E L O | 2 $|E_{LO}|^2$ is just a large, constant DC offset from the strong local oscillator.

The key term is the third one, the interference product: 2 E s E L O cos (…) $2\, E_{\mathrm{s}}\, E_{LO} \cos(\dots)$. This term has two magical properties. First, the weak signal field, E s E_{s}, is multiplied by the strong local oscillator field, E L O E_{LO}. This provides a huge amplification factor. We are effectively amplifying our signal optically, before it even hits the noise floor of the detector.

Second, this amplified signal now decays not with the time constant τ τ, but with a time constant of 2 τ 2τ. This is because the intensity is proportional to the field squared, but here we are detecting the field itself, E_s. The decay is slower, which can also improve detection sensitivity.

Page 58: The final piece of the heterodyne puzzle is demodulation

The key interference term that contains our amplified signal oscillates at the beat frequency, $\delta\nu$ (which was 40 kHz in our example schematic). All the other noise sources—detector noise, laser noise, etc.—are typically at different frequencies, often concentrated at low frequencies (DC).

By using a lock-in amplifier or a mixer to demodulate the detector’s output signal specifically at the frequency $\delta\nu$, we can selectively isolate and measure only the key interference term. This process effectively rejects all the noise that is not at the modulation frequency.

The result is a suppression of intensity noise and other technical noise sources by many orders of magnitude. This is what allows heterodyne-detected CRDS to achieve some of the highest sensitivities ever reported for absorption measurements, pushing deep into the shot-noise limit of detection.

Page 59

Let's discuss another important variant of CRDS, known as Cavity Leak-Out Spectroscopy, or CALOS.

CALOS is a variation of the CW-CRDS technique we discussed earlier. It uses continuous-wave lasers, but instead of just switching the beam off, it often involves sweeping either the laser frequency or the cavity length.

The procedure is as follows, in three steps:

1. The frequency of a CW laser is scanned over one of the cavity's resonances. 2. As the laser frequency sweeps into resonance, the intra-cavity power begins to build up dramatically. The transmitted light signal seen by the detector will trace out the Lorentzian profile of the cavity resonance. 3. When the transmitted power reaches a predetermined threshold level (usually near the peak of the resonance), a trigger is sent to rapidly block the input beam with an AOM. At that instant, the stored energy that has built up inside the cavity begins to "leak out," and the resulting ring-down decay is recorded, just as in other CRDS methods.

This trigger-and-block method allows for robust, repetitive measurements of the decay time.

Page 60:

CALOS offers several benefits over the standard pulsed CRDS method.

First, because it uses CW laser sources, it benefits from their inherently lower noise compared to many pulsed systems. Second, the alignment can be simpler. You don't need to handle the very high peak powers associated with short-pulse lasers, which reduces the risk of damaging the sensitive mirror coatings or other optical components. Finally, these advantages have led to some extremely high reported sensitivities. The slide notes a demonstrated sensitivity for $\alpha \leq 7 \times 10 -$

11 c m − 1 / H z $\alpha \leq 7 \times 10^{-11}\,\mathrm{cm}^{-1}/\sqrt{\mathrm{Hz}}$. This is truly state-of-the-art performance.

In essence, CALOS and its variants combine the high resolution and low noise of CW lasers with the fundamental time-domain measurement principle of CRDS.

Page 61:

This schematic nicely illustrates the workings of a Cavity Leak-Out Spectroscopy (CALOS) setup, specifically one that uses a swept cavity.

At the top left, we have a "cw Laser," which is often "Fibre-Coupled" for stability and ease of alignment. The light passes through "Mode-Matching Optics" and then an "AOM" which acts as a fast shutter. The light enters the "High-Finesse Optical Cavity" defined by mirrors M1 and M2.

Notice that Mirror M2 is mounted on a "PZT," or piezoelectric actuator. This PZT is driven by a "Ramp Generator" (Ramp Gen), which applies a sawtooth voltage to sweep the PZT, thereby sweeping the cavity length.

The light leaking out of M2 hits a "Photodetector." The signal goes to an "Oscilloscope / DAQ" for data acquisition.

Now look at the inset graph of the "Detected Signal Profile." As the ramp generator sweeps the cavity length, the cavity resonance sweeps past the fixed laser frequency. We see the signal "Build-up" as it comes into resonance. When the intensity hits a pre-set "Trigger Threshold," two things happen. A trigger is sent to the "AOM Driver" to switch off the laser beam (the RF switch). Another trigger is sent to the oscilloscope to start recording the decay. The moment the laser is blocked, we see the signal begin its exponential "Ring-down," which is the decay we measure. This entire cycle repeats with every sweep of the ramp generator.

Page 62:

Let's look at some of the practical achievements of the CALOS technique. Its high sensitivity has opened up a range of important applications.

In medical diagnostics, for example, CALOS is used for breath analysis. By analyzing trace amounts of certain molecules in exhaled breath, it's possible to detect disease markers. The slide mentions nitric oxide (NO), carbon monoxide (CO), ammonia (NH_3), and carbon dioxide (CO_2). For these species, detection limits can reach the parts-per-trillion concentration level, which is an extraordinary feat.

In atmospheric science, CALOS is used for real-time measurements of highly reactive and short-lived atmospheric radicals, such as the hydroxyl radical (OH) and the nitrate radical (NO_3). These species play a crucial role in atmospheric chemistry, and measuring their concentration with high temporal resolution is vital for understanding processes like smog formation and ozone depletion.

The advantages can be summarized as follows: CW lasers offer sub-megahertz linewidths, meaning the spectral resolution is almost always limited by the absorber itself (e.g., by Doppler broadening), not the instrument. Also, the ability to rapidly scan the laser or cavity allows for multiplexed detection, where the concentrations of several different molecular species can be monitored simultaneously or in rapid succession.

Page 63:

As with any real-world technique, CALOS has its limiting factors.

The primary one, which is common to all CRDS variants, is the availability of suitable high-reflectivity mirrors at the specific wavelengths of interest. This remains a key technological challenge, especially for moving further into the UV or mid-IR spectral regions.

For applications involving flowing gas samples, any residual absorption from the windows or walls of the sample flow cell can contribute to the baseline loss and must be carefully accounted for.

Finally, for swept-cavity implementations like the one we saw, the mechanical stability of the swept cavity is paramount. Any vibrations or drift in the PZT scanner can introduce noise and limit the ultimate precision of the measurement. These are engineering challenges that must be overcome to realize the full potential of the technique.

Page 64:

This slide describes a modern implementation that combines several of the advanced concepts we've discussed, often referred to as Rapidly-Swept CW-CRDS.

These state-of-the-art systems combine the best of all worlds: they use very high-reflectivity mirrors, they are often built using robust fiber-optic components, and they can be designed as a single-ended transmitter-receiver package, which is ideal for remote sensing applications like monitoring atmospheric pollutants over a long open path.

The slide gives a stunning performance example for the detection of acetylene gas, C 2 H 2 C_2H_2, at a wavelength of 1.525 μ m $1.525\ \mu$m, which is in the near-infrared telecommunications band where high-quality fiber components are readily available.

At low pressure, where the absorption line is Doppler-limited, a detection limit of just 19 m T o r r $19\ \mathrm{mTorr}$ is achieved. This corresponds to an astonishingly low partial pressure of 2.5 × 10 − 11 2.5×10^{-11} atmospheres.

To put this in a more practical context, this sensitivity allows for an ambient-pressure measurement of C 2 H 2 C_2H_2 down to a concentration of 0.37 0.37 parts-per-billion by volume, or ppbv. This is the level of sensitivity required for cutting-edge environmental monitoring and industrial process control.

Page 65:

The remarkable performance of these modern systems is enabled by several hardware optimizations.

These include the design of specialized low-absorption flow cells that minimize any unwanted background signals.

They rely on high-speed data acquisition systems that can rapidly digitize the decay traces and perform real-time fitting and averaging.

And they use precise, computer-controlled actuation of the cavity length, typically with piezoelectric transducers, to enable the rapid and stable scanning and locking that is required for these advanced CW techniques.

It is this combination of fundamental optical principles with sophisticated engineering and control that pushes the boundaries of spectroscopic detection.

Page 66:

Let's explore yet another clever variation in the CRDS family: Phase-Shift CRDS. This technique offers an alternative way to measure the ring-down time, not in the time domain, but in the frequency domain.

This method uses a continuous-wave laser, but its intensity is sinusoidally modulated at a specific angular frequency, which we'll call Ω $\mathit{\Omega}$.

The modulated laser beam is then split into two paths. One path, the "probe," is sent through the optical cavity. The other path, the "reference," bypasses the cavity.

Now, here is the key concept. The optical cavity acts like a low-pass filter; it has a finite response time characterized by its decay time, τ τ. When the intensity-modulated probe beam passes through this "slow" cavity, the modulation envelope itself acquires a phase lag, which we'll call by the Greek letter ϕ ϕ, relative to the reference beam's modulation.

There is a very simple and direct mathematical relationship between this measured phase lag, ϕ ϕ, the modulation frequency, Ω $\mathit{\Omega}$, and the cavity decay time, τ τ. The relationship is: the tangent of ϕ ϕ equals Ω $\mathit{\Omega}$ times τ τ.

$\tan(\phi) = \Omega \tau$.

$$\tan(\phi) = \mathit{\Omega}\,\tau.$$

So, by simply measuring a phase shift, we can directly calculate the ring-down time.

Page 67:

So, in a Phase-Shift CRDS experiment, the measurement process is as follows:

- We measure the phase shift, ϕ, as a function of the laser wavelength, λ, as we tune it across an absorption feature. - From the relationship $\tan(\phi) = \Omega\tau$, we can calculate τ as a function of λ. - And once we have $\tau(\lambda)$, we can use our standard CRDS formula to calculate the absorption coefficient, $\alpha(\lambda)$, and thus obtain the absorption spectrum.

It's interesting to note, as the final bullet point mentions, that this exact same principle is widely used in a different field of spectroscopy: fluorescence lifetime measurements. The technique known as "phase fluorometry" measures the phase lag of a modulated fluorescence signal relative to the modulated excitation light to determine the fluorescence lifetime of a molecule. The underlying physics of a linear system's response is identical. It's a beautiful example of how the same fundamental concepts appear in different contexts in physics.

Page 68:

We now come to our final, and perhaps one of the most powerful, variations of the technique: Fourier-Transform Cavity Ring-Down Spectroscopy, or FT-CRDS.

This method combines the high sensitivity of CRDS with the multiplex advantage of Fourier-transform spectroscopy. The key modification to the experimental setup is to place a Michelson interferometer *after* the ring-down cavity, just before the detector.

The primary advantage of this approach is that it allows you to simultaneously record the entire absorption spectrum within the bandwidth of your laser source. In a traditional CRDS experiment, you have to scan the wavelength step-by-step, measuring one spectral point at a time. In FT-CRDS, you get all the points at once. This is known as Fellgett's multiplex advantage, and it can provide a massive improvement in the signal-to-noise ratio for a fixed total acquisition time.

The experimental steps are as follows.

Page 69: Let's walk through the steps of an FT-CRDS measurement.

1. A ring-down decay event is initiated in the cavity, for example by injecting a broadband laser pulse. The decaying light pulse leaks out of the cavity.

2. This leaking light, which contains absorption information from a wide range of wavelengths, is directed into a Michelson interferometer. As the path-difference in the interferometer is scanned by moving one of its mirrors, the different spectral components within the beam are modulated at different frequencies.

3. A single detector at the output of the interferometer records the total intensity as a function of the interferometer's path difference. This recorded signal is called an interferogram. The computer then performs a numerical Fourier transform on this interferogram to retrieve the absorption spectrum, α α as a function of frequency, ν ν.

This technique is especially powerful when combined with modern broadband light sources, such as mid-infrared frequency combs. A frequency comb is like having tens of thousands of perfectly stable, narrow-linewidth CW lasers all at once. Combining a comb with FT-CRDS allows for the simultaneous, ultra-sensitive measurement of complex spectra across a vast spectral range, which is revolutionizing fields like molecular spectroscopy and trace gas analysis.

Page 70:

This schematic illustrates the concept of Fourier-Transform CRDS. It's a four-step process.

Step 1: The "Optical Cavity". This is our standard high-finesse cavity where the ring-down decay occurs. The light leaking out contains the absorption information.

Step 2: The "Michelson Interferometer". The ring-down decay light enters the interferometer. It's split by a beam splitter. One path goes to a fixed mirror, the other to a movable mirror. The two paths recombine, and the interference between them depends on the path difference introduced by the movable mirror.

Step 3: The "Detector". A single detector measures the intensity of the recombined beam, recording the interferogram as the movable mirror is scanned.

Step 4: "Computer & FFT". The recorded interferogram is sent to a computer, which performs a Fast Fourier Transform (FFT). The result of the FFT is the retrieved spectrum, shown as a plot of Absorption, α (ν) $\alpha(\nu)$, versus Frequency, ν ν. This allows us to see the absorption features of our sample.

Page 71:

We're now ready to offer some concluding remarks on the rich and varied family of CRDS techniques.

The first point is that the CRDS family—which includes the pulsed, CW, heterodyne, CALOS, phase-shift, and Fourier-transform variations we've discussed—provides an incredibly versatile and powerful toolbox for performing ultra-sensitive absorption spectroscopy. These techniques are applicable across a wide range of wavelengths, from the ultraviolet to the infrared.

The second point is that for any of these techniques, the design and performance are governed by a few core parameters. These are: the mirror reflectivity, R R, which is the most critical parameter for determining the ultimate sensitivity; the cavity length, L L; the laser linewidth, which determines the spectral resolution; and the specific detection scheme chosen, which affects the noise characteristics and complexity of the system. A careful optimization of these parameters is key to building a successful experiment.

Page 72:

A crucial takeaway message is that the limitations of CRDS are primarily technical, rather than being rooted in fundamental physics. The underlying principles are sound. The main challenges are engineering ones: how to manufacture mirrors with even higher reflectivity and lower loss across broader wavelength ranges, and how to design more stable and robust mechanical and optical systems for alignment and locking.

The field is by no means static. Ongoing research continues to push the boundaries of what is possible. Exciting new developments include the integration of CRDS with optical frequency combs to provide broadband, high-resolution spectra with unprecedented speed and accuracy. There is also significant work in developing chip-scale cavities and integrating systems with fiber lasers to create compact, portable, and robust sensors.

These advancements are enabling the deployment of CRDS in a wide array of important field applications, including real-time environmental monitoring, noninvasive biomedical diagnostics through breath analysis, and even astrochemistry, where it can be used to study the composition of interstellar clouds in the laboratory.

CRDS began as a clever lab technique, but it has matured into a cornerstone of modern optical sensing and spectroscopy. That concludes our lecture for today. Thank you.