Chapter 5.6

Page 1:

Alright everyone, welcome back to Physics 608, Laser Spectroscopy. Today, we embark on a fascinating and fundamentally important topic: Chapter 5, Section 6, which deals with the **Linewidths of Single-Mode Lasers**. This material has been prepared by Distinguished Professor Doctor M A Gondal.

Understanding laser linewidth is not just an academic exercise; it's absolutely crucial for anyone working at the forefront of laser applications, especially in spectroscopy, metrology, and quantum technologies. We often think of lasers as sources of perfectly monochromatic light, but as we'll see, reality is more nuanced. We're going to delve into what determines the ultimate sharpness of a laser's emission, the fundamental limits imposed by quantum mechanics, and how these limits compare to what we can achieve in the lab. So, let's begin.

Page 2:

Let's start with the core motivation: **Why Does a Laser Have a Linewidth at All?** It’s a very pertinent question. After all, the idealized picture of a laser is a source of perfectly single-frequency light.

The first bullet point sets the stage:

Every real laser emits a spectrum with non-zero width, called the linewidth, denoted as capital Delta nu sub L ($\Delta\nu_{L}$).

This might seem counterintuitive at first. If a laser operates on a single mode of an optical resonator, shouldn't that mode correspond to an infinitesimally sharp frequency? Well, as we'll explore, there are inherent processes, even in an ideal single-mode laser, that cause the phase of the emitted light to fluctuate randomly over time. This phase fluctuation is what ultimately leads to a broadening of the spectral line, giving it a finite width, $\Delta\nu_{L}$. This $\Delta\nu_{L}$ is typically defined as the Full Width at Half Maximum, or FWHM, of the laser's power spectrum.

Now, why should we care about making this $\Delta\nu_L$ as small as possible? The second bullet point highlights this:

Ultranuarrow linewidths are essential for a variety of cutting-edge applications.

Let's look at some examples:

1. **High-resolution spectroscopy.** This is, of course, directly relevant to our course. If you want to resolve very fine details in the spectrum of an atom or molecule – for instance, **resolving isotope shifts**, which are tiny differences in transition frequencies between different isotopes of the same element, or **hyperfine splittings**, which arise from the interaction of nuclear spin with the electron cloud – you absolutely need a laser whose own linewidth is significantly narrower than the features you're trying to observe. If your laser linewidth is broader than the separation of these subtle spectral features, those features will simply be washed out, convolved with your laser line, and you won't be able to distinguish them. So, for high-fidelity spectroscopic interrogation, an ultranarrow laser is your precision scalpel.

Page 3:

Continuing with the applications where ultranarrow linewidths are indispensable:

2. **Optical frequency standards.** This is a domain where lasers have revolutionized precision. Think of **atomic clocks**. The most advanced atomic clocks today use optical transitions in atoms as their "pendulum." The stability and accuracy of these clocks are directly tied to the ability to lock a laser with an extremely narrow linewidth to these atomic transitions. The narrower the laser line, the more precisely its frequency can be determined and stabilized, leading to clocks with astounding accuracy. Another area is **laser interferometric length metrology**. For measuring distances with extreme precision using interference, you need a light

source with a very long coherence length, which, as we'll see, is inversely related to the linewidth. The narrower the linewidth, the longer the distance over which the phase of the light remains predictable, allowing for highly precise interferometric measurements.

3. **Coherent manipulation of quantum states in quantum information.** In fields like quantum computing or quantum communication, we often need to use laser light to precisely control the quantum state of atoms, ions, or other quantum bits (qubits). These manipulations rely on maintaining well-defined phase relationships. A broad laser linewidth implies rapid phase fluctuations, which can destroy the delicate quantum coherence needed for these operations. Therefore, lasers with extremely narrow linewidths are critical tools for encoding, manipulating, and reading out quantum information.

Now, beyond these practical applications, there's a deep fundamental physics aspect to understanding laser linewidths:

* **Understanding the ultimate, fundamental lower limit on** $\Delta\nu_L$ reveals deep connections between several core concepts in physics. This isn't just about engineering better lasers; it's about understanding the interplay of light and matter at the quantum level. Specifically, it connects to:

1. **Quantum fluctuations of the electromagnetic field.** Even in a perfect vacuum, the electromagnetic field is not quiescent. It has zero-point energy and undergoes fluctuations. These vacuum fluctuations can interact with the lasing process and are a source of fundamental noise, contributing to the laser linewidth.

2. **Spontaneous emission of atoms.** While stimulated emission is the heart of laser action, atoms in the gain medium can also emit photons spontaneously. These spontaneously emitted photons have random phases and directions, and those that get coupled into the lasing mode can perturb the phase of the coherent laser field, again contributing to the linewidth.

Page 4:

And the third deep connection that understanding the fundamental laser linewidth limit reveals is:

3. **Cavity photon statistics and phase diffusion.** The light field inside the laser cavity consists of a large number of photons. The process of spontaneous emission adding photons with random phases to this existing field leads to what's called phase diffusion – the phase of the laser's electromagnetic field undergoes a random walk over time. The statistics of the photons within the cavity, and how these random phase kicks accumulate, are central to determining the fundamental linewidth.

So, with this motivation and these connections in mind, we can state the **Goal of this section**:

- Our primary objective is to **derive the Schawlow-Townes quantum limit.** This is a famous result that predicts the minimum possible linewidth a laser can theoretically achieve, based on fundamental quantum principles.

- Alongside the derivation, we will **analyse the underlying noise mechanisms** – spontaneous emission, vacuum fluctuations, and their manifestation as phase diffusion – that give rise to this limit.

- And finally, we will **compare this theoretical quantum limit with experimentally achievable linewidths** in real-world lasers. This will give us a sense of how close we are to fundamental limits and what challenges remain.

This journey into the Schawlow-Townes limit is a classic topic in laser physics, and it beautifully illustrates the interplay of quantum mechanics, electromagnetism, and statistical physics in determining a key characteristic of lasers.

Page 5:

To visualize what we're talking about, this page presents a **Laser Linewidth Comparison** graph.

Let's break down this illustration. On the vertical axis, we have **Relative Intensity**. On the horizontal axis, we have **Frequency, nu (**ν). The graph shows three idealized spectral profiles of a laser line, all centered around the same nominal frequency, but with vastly different widths.

1. The outermost, broadest curve, shaded in **orange**, is labeled **Practical**. This represents the linewidth of a typical, everyday laser that you might find in a teaching lab or a common industrial application. While it's still narrow compared to, say, an incandescent bulb, its linewidth is significantly broadened by technical noise sources – vibrations, temperature fluctuations, power supply instabilities, and so on. We're talking about linewidths that could be in the megahertz or kilohertz range, or sometimes even broader for less sophisticated lasers.

2. The middle curve, shaded in a **greenish-blue (cyan)**, is labeled **State-of-the-art**. This represents a highly engineered laser system where significant effort has been made to suppress technical noise. This could involve sophisticated thermal and vibrational isolation, ultra-stable power supplies, and feedback control systems. These lasers, often found in research labs pursuing precision measurements, can achieve linewidths in the Hertz to kiloHertz range. This is a significant improvement over the "practical" laser.

3. Finally, the innermost, sharpest peak, shaded in **blue**, is labeled **Quantum-limited**. This represents the ultimate theoretical limit to the laser linewidth, as predicted by the Schawlow-Townes formula (and its refinements). This linewidth arises purely from fundamental quantum noise – spontaneous emission and vacuum fluctuations – assuming all technical noise has been perfectly eliminated. As you can see, this quantum-limited linewidth is *extremely* narrow, often in the milliHertz or even microHertz range for typical laser parameters.

The key takeaway from this visual is the vast difference in scales between what's practically common, what's achievable with advanced engineering, and what fundamental quantum mechanics dictates as the ultimate limit. Our goal in the upcoming discussion is to understand the physics behind that innermost blue curve – the quantum limit.

Page 6:

Now, let's move on to some **Basic Definitions**, as we transition **From Perfect Monochromaticity to Finite Bandwidth**.

First, consider the ideal scenario:

* An **Ideal monochromatic wave**. This is a purely theoretical construct, **never realised in practice.**

Mathematically, we would write the electric field of such a wave as:

$$E_{\text{ideal}}(t) = E_0 \cos(2\pi\nu_0\, t)$$

That is, capital E subscript ideal, as a function of time t, equals capital E subscript zero, times cosine, of the quantity two pi nu subscript zero t.

Here, **capital E sub zero (E_0)** is the constant amplitude of the wave, representing its strength, typically in Volts per meter. And **nu sub zero (v_0)** is the perfectly defined, single frequency of the wave, in Hertz. In this ideal case, if you were to take a Fourier transform to get the spectrum, you would find an infinitely sharp delta function at the frequency ν_0. The phase is perfectly predictable for all time.

Why is this never realized in practice? Because any real oscillator, including a laser, is subject to perturbations and has a finite energy or finite observation time, which inherently leads to some frequency spread due to the uncertainty principle, and more practically, due to the noise processes we've started to discuss.

Now, let's consider a more realistic description:

* A **Real laser field must be written** in a way that accounts for these imperfections.

The electric field of a real laser is typically expressed as:

$$E(t) = A(t)\cos\big(2\pi\nu_0\, t + \phi(t)\big)$$

That is, capital E as a function of time t, equals capital A as a function of time t, times cosine, of the quantity two pi nu sub zero t, plus phi as a function of time t.

Here, **nu sub zero (ν_0)** is still the nominal center frequency of the laser. However, we now have two new time-dependent terms: **Capital A of t (A(t))** represents the amplitude of the laser field, which can fluctuate over time. And, crucially for our discussion of linewidth, **phi of t (φ(t))** represents the phase of the laser field, which can also fluctuate over time. It's these fluctuations, particularly in $\phi(t)$, that are primarily responsible for the laser's finite linewidth.

Page 7:

Let's elaborate on the terms in our expression for the real laser field, $E(t) = A(t)\cos[2\pi\nu_0 t + \phi(t)]$.

* $A(t)$ is the slowly varying amplitude, typically measured in **Volts per meter (**$\mathrm{V\,m^{-1}}$). The term "slowly varying" means that $A(t)$ changes on a timescale that is much longer than the period of the optical oscillation ($1/\nu_0$). Fluctuations in $A(t)$ are referred to as amplitude noise.

* $\phi(t)$ is the slowly varying phase, measured in **radians (**rad). Similarly, "slowly varying" means $\phi(t)$ changes on a timescale much longer than the optical period. Fluctuations in $\phi(t)$ are referred to as phase noise or frequency noise (since instantaneous frequency is related to the time derivative of the total phase). As we'll see, phase noise is the dominant contributor to the fundamental laser linewidth.

Now, let's define the **Linewidth, capital Delta nu sub L (**$\Delta\nu_{\text{L}}$), more formally:

* It is the **Full Width at Half Maximum (FWHM)** of the **power spectral density** of the laser field. The power spectral density is obtained by taking the Fourier transform of the electric field $E(t)$ to get $E(\nu)$ (E as a function of frequency ν), and then considering its magnitude squared, **absolute value of E of nu, squared (**$|E(\nu)|^2$). This $|E(\nu)|^2$ represents how the laser's power is distributed across different frequencies. The FWHM is the width of this spectral peak at the points where the power has dropped to half of its maximum value.

This linewidth, $\Delta\nu_{\text{L}}$, is intimately related to another important concept: the **coherence time, tau sub c (**τ_{c}). The relationship is approximately:

tau sub c is approximately equal to one divided by (pi times capital Delta nu sub L) That is, $\tau_{\text{c}} \approx \frac{1}{\pi\Delta\nu_{\text{L}}}$.

The coherence time, τ_{c}, represents the average time interval over which the phase of the laser field remains predictable or correlated with itself. A narrower linewidth $\Delta\nu_{\text{L}}$ implies a longer coherence time τ_{c}. The factor of π arises from the Fourier transform relationship between a Lorentzian lineshape (which is often a good model for a laser line broadened by phase diffusion) and an exponential decay of its time-domain correlation function.

Building on coherence time, we have:

* **Coherence length, capital L sub c (**L_{c}): This is the **distance over which the field remains phase-correlated.** It's simply the distance light travels in one coherence time.

Page 8

The relationship between coherence length L_{c}, the speed of light c, and coherence time τ_{c} is straightforward:

$$L_{\text{c}} = c\tau_{\text{c}}$$

That is, $L_{\text{c}} = c\tau_{\text{c}}$.

Here, c is the speed of light in the medium (usually vacuum or air for laser propagation). So, if you know the coherence time, you can easily find the coherence length.

The implication of this is very important:

Narrower capital Delta nu sub L ($\Delta\nu_{\text{L}}$) leads to a longer capital L sub c (L_{c}).

This is a direct consequence of the relationships we just discussed: narrower $\Delta\nu_{\text{L}}$ means longer τ_{c}, and longer τ_{c} means longer L_{c}.

To give you a sense of scale, the slide notes: **(kilometres for Hertz-level lasers).**

Think about that. If you have a laser with a linewidth of just 1 Hertz ($\Delta\nu_{\text{L}} = 1\,\text{Hz}$), then its coherence time τ_{c} would be approximately

$$\tau_{\text{c}} \approx \frac{1}{\pi \cdot 1\,\text{Hz}} \approx 0.3\,\text{seconds}$$

The coherence length L_{c} would then be

$$L_{\text{c}} = c\,\tau_{\text{c}} \approx (3 \times 10^{8}\,\text{m/s}) \times 0.3\,\text{s} \approx 0.9 \times 10^{8}\,\text{m}$$

which is 90,000 kilometers! This is an incredibly long distance, highlighting the remarkable phase stability of such lasers. This long coherence length is precisely what enables applications like long-baseline interferometry for gravitational wave detection, or high-precision metrology over significant distances.

So, the quest for narrower linewidths is also a quest for longer coherence times and coherence lengths, opening up new possibilities for precision measurement and coherent control.

Page 9:

Now, let's delve into the **Classification of Frequency Noise in Lasers**. Understanding the sources of noise is crucial for understanding linewidth and for devising strategies to reduce it. We can broadly categorize noise into two types:

First, there's **Technical (or extrinsic) noise**:

* This type of noise **originates outside quantum theory.** It's due to classical, often environmental, perturbations to the laser system. These are, in principle, things we can try to engineer away, though in practice, it's very challenging to eliminate them completely. Examples include: * **Fluctuations of cavity length,** d. The resonant frequencies of a laser cavity are determined by its length. If the cavity length changes due to **thermal expansion** (e.g., the laser heats up or the room temperature changes) or **vibrations** (mechanical disturbances from the environment like acoustic noise, pumps, or even seismic activity), the resonant frequencies will shift, leading to frequency noise. * **Variations of refractive index,** n, of intracavity components. If there are materials inside the laser cavity (like the gain medium itself, or optical elements), their refractive index can change due to **temperature fluctuations, pressure changes, or even acoustic waves** passing through them. Changes in refractive index alter the optical path length of the cavity, which, similar to physical length changes, shifts the resonant frequencies. * **Power-supply and pump-intensity noise.** Fluctuations in the electrical power supplied to the laser, or instabilities in the intensity of the pump source (e.g., a flashlamp or another laser used to energize the gain medium), can lead to variations in the gain, temperature, or carrier density within the laser, all of which can translate into frequency noise.

Second, and more fundamentally, there's **Fundamental (or intrinsic) noise**:

* This type of noise **cannot be eliminated even with "perfect" engineering** because it arises from the quantum nature of light and matter. It sets the ultimate lower limit on the laser linewidth, which is what the Schawlow-Townes theory addresses. The primary sources of fundamental noise are: 1. **Spontaneous emission into the lasing mode.** Even in a laser operating well above threshold, atoms in the gain medium will spontaneously emit photons. While most of these go off in random directions, a small fraction will be emitted into the same spatial and polarization mode as the laser field. Each such spontaneously emitted photon has a random phase relative to the existing coherent field in the mode. These random phase kicks perturb the overall phase of the laser light, causing it to diffuse over time. This is a key mechanism. 2. **Photon-number (shot) noise leading to amplitude fluctuations.** The emission of photons is a discrete, quantum process. Even for a perfectly stable average intensity, there will be statistical fluctuations in the number of photons emitted per unit time. This is known as shot noise. While primarily affecting the amplitude, these amplitude fluctuations can also couple, to a lesser extent, into phase fluctuations through nonlinearities in the gain medium (like the Kramers-Kronig relations linking gain and refractive index).

Page 10:

Continuing with the sources of fundamental (intrinsic) noise:

3. **Phase diffusion generated by random birth times of photons.** This is essentially a restatement and consequence of point 1 (spontaneous emission). Each time a spontaneous photon is added to the lasing mode, it carries a random phase. The accumulation of these random phase "kicks" causes the overall phase of the laser field to undergo a random walk, a process known as phase diffusion. This continuous, random evolution of the phase is the direct cause of the fundamental linewidth. Think of it like a drunken sailor's walk – each step is random, and over time, the sailor can end up quite far from the starting point, even if each individual step is small. Similarly, the laser's phase drifts.

Now, an important clarification for the subsequent discussion:

* The **Linewidth discussion below completely ignores technical noise and addresses only the intrinsic limit.**

We are going to focus on deriving the Schawlow-Townes limit, which is a quantum limit. We will assume that our hypothetical laser is perfectly engineered, meaning all sources of technical noise (vibrations, temperature drifts, etc.) have been eliminated. This allows us to isolate and understand the fundamental noise processes that set the ultimate floor on the laser linewidth. In reality, achieving this intrinsic limit is extremely difficult, as technical noise often dominates. But knowing the fundamental limit is crucial as it provides a benchmark and a target for experimental efforts.

Page 11

This page presents a **Schematic Comparison of Noise Sources in a Well-Stabilized Single-Mode Laser**, in the form of a bar chart. This visual helps to put the magnitudes of different noise sources into perspective.

On the vertical axis, we have **Relative Noise Magnitude**, on an arbitrary linear scale, let's say from 0.0 to 3.0. The horizontal axis categorizes different noise sources.

The sources are grouped into two categories:

1. **Technical Noise Sources** (shown with blue bars):

Cavity ΔL (Cavity length fluctuations): This bar is shown as the largest, with a relative magnitude of about 2.5. This signifies that, even in a well-stabilized laser, mechanical stability of the cavity length is often a dominant challenge.

Index Δn (Refractive index fluctuations): This bar is next, with a magnitude around 2.0. Changes in the refractive index of intracavity materials due to temperature or pressure are also significant.

Power/Pump (Power supply or pump fluctuations): This bar is somewhat smaller, around 1.5, but still a considerable source of technical noise.

2. **Fundamental Noise Sources** (shown with red bars):

Spont. Emission (Spontaneous Emission into the lasing mode): This is the largest of the fundamental noise sources, with a magnitude of about 1.0. This is the primary driver of phase diffusion that leads to the Schawlow-Townes linewidth.

Phase Diffusion: This bar is shown with a magnitude of about 0.8. It's essentially the consequence of spontaneous emission, representing the random walk of the laser's phase.

Shot Noise: This is the smallest, with a magnitude around 0.5. This refers to the quantum fluctuations in photon number, which primarily affect amplitude but can also have a minor coupling to phase.

Important context: The title specifies "in a Well-Stabilized Single-Mode Laser." This is key. If the laser were *not* well-stabilized, the blue bars representing technical noise would be vastly taller, completely dwarfing the red bars of fundamental noise. The chart illustrates that even after significant engineering efforts to suppress technical noise, these technical sources can still be comparable to or even larger than the fundamental quantum noise sources. Reaching the quantum limit requires heroic efforts to reduce those blue bars down to the level of, or below, the red bars.

This chart effectively summarizes the battle faced by experimentalists: first, to aggressively combat the large technical noise sources, and then, to confront the unavoidable fundamental quantum noise. Our theoretical discussion will focus on the origins and magnitude of those red bars.

Page 12:

Let's now focus on the **First Fundamental Noise Source: Spontaneous Emission**. This is a cornerstone of the Schawlow-Townes theory.

The first bullet point describes the basic process:

* An **Atom in the upper laser level,** E_i (Eᵢ), decays to a lower level, E_k (E□), with a probability per unit time denoted as A_{ik} (Aᵢ□). This A_{ik} is the **Einstein A-coefficient** for spontaneous emission between levels i and k, and it has units of **inverse seconds (**s^{-1}). So, A_{ik} represents the rate at which an isolated atom in the excited state E_i will spontaneously jump down to state E_k, emitting a photon in the process. This emission occurs without any external stimulation, hence "spontaneous." The energy of the emitted photon will be $h\nu = E_i - E_k$.

Now, consider the collective effect in the laser's active medium:

* The **Total spontaneous power radiated from an active volume,** V_m (V□), containing a population density N_i (Nᵢ) in the upper laser level is given by the following equation. Let P_{sp} (Psp) be this total spontaneous power:

$$P_{sp} = N_i V_m A_{ik} h\nu_L$$

That is, capital P_{sp}, equals capital N_i, times capital V_m, times capital A_{ik}, times h (Planck's constant) times ν_L (the laser frequency).

Let's break down the terms in this equation: * P_{sp} (Psp): Total power (energy per unit time, in Watts) emitted spontaneously from the active volume. * N_i (Nᵢ): The population density of atoms in the upper laser level E_i (number of atoms per unit volume, e.g., in units of cm⁻ ³ or m⁻ ³). * V_m (V□): The volume of the active gain medium (e.g., in cm³ or m³). So, N_{iV_m} is the total number of atoms in the upper laser level. * A_{ik} (Aᵢ□): The Einstein A-coefficient, or rate of spontaneous emission per atom (in s^{-1}). * h: Planck's constant (approximately 6.626×10^{-34} Joule-seconds). * ν_L (vL): The laser frequency (in Hertz), which is approximately $\frac{E_i - E_k}{h}$. So, $h\nu_L$ is the energy of a single spontaneously emitted photon.

The final bullet simply clarifies:

* h = Planck's constant, ν_{L} = laser frequency.

The crucial point for laser linewidth is that *some* of this spontaneously emitted power, with its random phase, will inevitably be directed into the specific mode in which the laser is oscillating.

Page 13:

Now, let's consider where these spontaneously emitted photons go.

- **Photons are emitted into *all* available electromagnetic modes within the fluorescence bandwidth (often characterized by the Doppler width,** $\Delta\nu_{\mathrm{D}}$, $\Delta\nu_{\mathrm{D}}$). This is a critical distinction from stimulated emission, which preferentially adds photons *into the existing lasing mode*. Spontaneous emission is, by its nature, random. An excited atom doesn't "know" about the laser field when it decides to emit spontaneously (unless it's undergoing stimulated emission instead). So, these photons can go into any direction, any polarization, and any frequency allowed by the energy difference and the lineshape of the transition. The fluorescence bandwidth, $\Delta\nu_{\mathrm{D}}$, (for example, the Doppler width in a gas laser, or a more general gain bandwidth in other lasers) defines the range of frequencies over which these spontaneous emissions occur.

- **The Number of spatial and polarisation modes per unit volume in a given bandwidth,** $\Delta\nu_{\mathrm{D}}$, is given by a well-known formula from electromagnetic theory. Let N_{modes} represent this density of modes: N_{modes} is approximately equal to $\frac{8\pi\nu_{\mathrm{L}}^2\Delta\nu_{\mathrm{D}}}{c^3}$

That is,

$$N_{\mathrm{modes}} \approx \frac{8\pi\nu_{\mathrm{L}}^2\Delta\nu_{\mathrm{D}}}{c^3}.$$

Let's examine the terms:

- N_{modes}: Number of electromagnetic modes per unit volume per unit frequency range (if $\Delta\nu_{\text{D}}$ were unity), but here it's the number of modes per unit volume *within the bandwidth* $\Delta\nu_{\text{D}}$. So its units would be something like modes per cubic meter.

- ν_{L}: The central laser frequency (in Hertz). The ν_{L}^2 term indicates that the density of modes increases rapidly with frequency.

- $\Delta\nu_{\text{D}}$: The fluorescence bandwidth (e.g., Doppler width) over which spontaneous emission occurs (in Hertz).

- c: The speed of light in the medium (in meters per second). The c^3 in the denominator comes from the three-dimensional nature of space and the dispersion relation for light.

- The factor of 8π includes a factor of 2 for two independent polarizations, and factors related to integrating over solid angle and frequency space.

This formula tells us that there is a vast number of electromagnetic modes available for spontaneously emitted photons to go into, especially at optical frequencies where ν_{L} is large. Only a tiny fraction of these will happen to align with the specific single mode of the laser resonator.

Page 14:

Building on the concept of many available modes for spontaneous emission, we arrive at a very important consequence:

* The **Mean photon number per mode from spontaneous emission is much, much less than** 1 ($\ll 1$).

This means that, if you pick any single one of those numerous electromagnetic modes (including the specific mode that will eventually lase), the average number of photons spontaneously emitted into *that particular mode* at any given time (before lasing action dominates) is very small, typically a tiny fraction of a single photon.

Therefore, the initial optical field is dominated by vacuum fluctuations until the laser reaches threshold.

Before the laser turns on and stimulated emission takes over to populate the lasing mode with a large number of coherent photons, that mode is essentially "empty" or, more accurately, contains only the zero-point energy associated with vacuum fluctuations. The occasional spontaneously emitted photon that happens to find its way into this mode is a rare event.

This has a profound implication for how a laser starts up: it's these vacuum fluctuations, or a "seed" photon from spontaneous emission, that initiate the process of stimulated emission. Once stimulated emission begins to build up a significant photon population in the cavity mode (i.e., the laser reaches threshold), then the field in that mode is no longer dominated by vacuum fluctuations but by the coherent, amplified laser light.

However, even above threshold, spontaneous emission continues. And those spontaneous photons that *do* get coupled into the lasing mode, though few in number compared to the stimulated photons, are the ones that carry random phases and contribute to the fundamental laser linewidth. The fact that the mean number *per mode* is small highlights that only a tiny fraction of the total spontaneous emission actually ends up in the lasing mode, but that tiny fraction is critically important for the linewidth.

Page 15:

This slide provides a helpful visual: **Spontaneous Emission: Energy Levels and Emission Bandwidth.** It consists of two diagrams.

On the left side, we see the "Spontaneous Emission Process" illustrated with an energy level diagram:

* The vertical axis represents **Energy**. * Two horizontal black lines denote two energy levels: an **upper laser level labeled** E_i and a **lower laser level labeled** E_k. * A **blue circle** on the level E_i represents an atom (or molecule, or ion) in the excited state. * A **black downward arrow** shows this atom

transitioning from E_i to E_k. This represents the decay process. * From the atom, several **red wavy arrows** emanate, symbolizing the spontaneously emitted photons. These arrows point in various directions, emphasizing the random nature of the emission direction. One of these red arrows is labeled $h\nu_\mathrm{L}$, indicating the energy of the emitted photon. * The transition itself is labeled with A_{ik}, the Einstein A-coefficient, representing the probability per unit time for this spontaneous decay.

This diagram visually reinforces that spontaneous emission from an excited atom can result in photons going into many different directions (modes). Only if one of these red arrows happens to align perfectly with the laser cavity's optical axis and mode structure will it contribute to the field within that specific lasing mode.

On the right side, we see a graph representing the "Emission Spectrum":

* The vertical axis is **Intensity**, and the horizontal axis is **frequency,** ν. * A **blue, bell-shaped curve** (typically Gaussian for Doppler broadening, or Lorentzian for natural/collisional broadening) is plotted. This curve represents the **fluorescence bandwidth** or the **gain bandwidth** of the transition. It shows the range of frequencies over which spontaneous emission can occur. * The spectrum is centered around the nominal laser frequency, labeled ν_L on the frequency axis (though shown as ν_L with a tick mark, it implies the center). * A horizontal double-arrow line segment near the half-maximum points of the curve is labeled $\Delta\nu_\mathrm{D}$. This indicates the **Full Width at Half Maximum (FWHM)** of the emission spectrum, which could be, for example, the Doppler width in a gas laser.

This spectrum shows that spontaneously emitted photons are not all at exactly the same frequency ν_L, but are spread over a range $\Delta\nu_\mathrm{D}$. If the laser cavity has a very narrow resonance (narrow $\Delta\nu_\mathrm{c}$), it will only effectively interact with or amplify those spontaneously emitted photons whose frequencies fall within that narrow cavity resonance.

Together, these two diagrams illustrate that spontaneous emission is a random process both in terms of direction and, to some extent, frequency (within $\Delta\nu_D$). This randomness is the root of its contribution to laser noise.

Page 16:

Let's now put some numbers to these concepts with a **Numerical Illustration – Spontaneous Photons in a Helium-Neon (HeNe) Laser.** The HeNe laser is a common, well-understood system, often used as an example. We're considering the red HeNe line at a wavelength $\lambda = 632.8$ nanometers.

Here are some typical parameters:

* **Stationary upper-level population density:** $N_i \approx 10^{10}\,\text{cm}^{-3}$. This N_i is the number of helium or neon atoms (specifically, the species responsible for the upper state of the lasing transition) per unit volume that are in the excited state ready to lase or emit spontaneously.

* **Einstein coefficient:** $A_{ik} \approx 10^8\,\text{s}^{-1}$. This means that an atom in the upper laser level will, on average, spontaneously emit a photon and decay to the lower level about 100 million times per second if left on its own.

From these, we can calculate the:

* **Spontaneous photon rate per unit volume:** Let this be R_{sp}. $R_{sp} = N_i A_{ik} = 10^{18}\,\text{photons}\,\text{s}^{-1}\,\text{cm}^{-3}$. This is a huge number! Every cubic centimeter of the active medium is spewing out 10^{18} spontaneous photons every second. These are radiated in all directions over the fluorescence bandwidth.

Now, let's consider the modes:

* **Modes in Doppler bandwidth at a wavelength** $\lambda = 632.8$ nanometers: We need to use the formula for N_{modes} from page 13:

$$N_{\text{modes}} \approx \frac{8\pi\, \nu_{\text{L}}^{2}\, \Delta\nu_{\text{D}}}{c^3}.$$

First, let's find ν_{L}:

$$\nu_{\text{L}} = \frac{c}{\lambda} = \frac{3 \times 10^{10}\,\text{cm/s}}{632.8 \times 10^{-7}\,\text{cm}} \approx 4.74 \times 10^{14}\,\text{Hz}.$$

For a HeNe laser, the Doppler width $\Delta\nu_{\text{D}}$ is typically around 1.5 GHz (1.5×10^{9} Hz). Plugging these into the formula for N_{modes}:

$$N_{\text{modes}} \approx \frac{8\pi \cdot (4.74 \times 10^{14}\,\text{Hz})^2 \cdot (1.5 \times 10^{9}\,\text{Hz})}{(3 \times 10^{10}\,\text{cm/s})^3}.$$

$$N_{\text{modes}} \approx \frac{8\pi \cdot 2.25 \times 10^{29}\,\text{Hz}^{2} \cdot 1.5 \times 10^{9}\,\text{Hz}}{2.7 \times 10^{31}\,\text{cm}^{3}\,\text{s}^{-3}}.$$

$$N_{\text{modes}} \approx \frac{8.48 \times 10^{48}\,\text{Hz}^{3}}{2.7 \times 10^{31}\,\text{cm}^{3}\,\text{s}^{-3}}.$$

Since $\text{Hz} = \text{s}^{-1}$, $\text{Hz}^3 = \text{s}^{-3}$, this cancels out.

$$N_{\text{modes}} \approx 3.14 \times 10^{17}\,\text{modes/cm}^{3}.$$

The next slide gives a value for N_{modes}, let's see how it compares. It seems the calculation here is for N_{modes} (density of modes).

Page 17:

Continuing our numerical illustration for the HeNe laser:

The slide states: $N_{\text{modes}} \approx 3 \times 10^{8}$ modes cm^{-3}.

There seems to be a discrepancy between this value and my quick calculation on the previous page, which yielded a much larger number (around 3×10^{17} modes cm^{-3}). The formula for N_{modes} is

$$N_{\text{modes}} = \frac{8\pi\nu L^2 \Delta\nu_{\text{D}}}{c^3}.$$

Perhaps the $\Delta\nu_{\text{D}}$ used for the slide's value is much smaller, or there's a different interpretation of N_{modes} here. For the sake of continuity with the slide's argument, let's proceed using the value provided on the slide, which is $N_{\text{modes}} \approx 3 \times 10^8$ modes cm^{-3}. This value represents the density of electromagnetic modes available within the Doppler broadened fluorescence bandwidth.

Now, let's calculate the **Photon flux into a single mode**:

The slide presents an equation:

$$\Phi = \frac{R_{\text{sp}}}{N_{\text{modes}}} \approx 3 \times 10^9 \text{ photons s}^{-1}.$$

Let's check this using the slide's R_{sp} and N_{modes}:

$$\Phi = \frac{10^{18} \text{ photons s}^{-1} \text{ cm}^{-3}}{3 \times 10^8 \text{ modes cm}^{-3}}$$

$$\Phi \approx 3.33 \times 10^9 \text{ photons s}^{-1} \text{ per mode.}$$

This result seems consistent. So, Φ here represents the rate at which spontaneous photons are emitted into *any single, specific mode* (like the lasing mode). While the total rate R_{sp} is enormous (10^{18} photons s^{-1} cm^{-3}), it's distributed among a very large number of available modes (3×10^8 modes cm^{-3} according to the slide). Thus, the rate into any *one* mode is substantially smaller, about 3 billion photons per second per mode.

Next, the **Mean photon density in that mode**:

The slide gives an equation:

$$\langle n_{\text{ph}} \rangle = \frac{\Phi}{c} \leq 0.1.$$

Here, $\langle n_{\text{ph}} \rangle$ represents the average number of spontaneously emitted photons *occupying* that single lasing mode at any given instant.

Let's analyze Φ/c: Φ is in photons s^{-1} per mode. c is the speed of light (3×10^{10} cm s^{-1}). So Φ/c would have units of

$$\frac{\text{photons s}^{-1}\text{ mode}^{-1}}{\text{cm s}^{-1}} = \text{photons cm}^{-1}\text{ mode}^{-1},$$

which is a linear density of photons per mode. However, the result "≤ 0.1" is dimensionless, suggesting $\langle n_{\text{ph}} \rangle$ is an average *number* of photons, not a density.

This expression $\langle n_{\text{ph}} \rangle = \frac{\Phi}{c}$ is a common simplification, often meaning $\Phi \times \tau_{\text{photon}}$ where τ_{photon} is a characteristic time like $1/c$ multiplied by a characteristic length of the mode, or more accurately, $\Phi \times \tau_{\text{cavity}}$ where τ_{cavity} is the photon lifetime in the cavity for that mode.

If τ_{cavity} for a HeNe laser is, for instance, around 30 picoseconds (3×10^{-11} s), then

$$\begin{aligned} \langle n_{\text{ph}} \rangle &\approx \left(3 \times 10^{9}\text{ photons s}^{-1}\text{ mode}^{-1}\right) \times (3 \times 10^{-11}\text{ s}) \\ &\approx 0.09\text{ photons per mode.} \end{aligned}$$

This is consistent with the "≤ 0.1" value.

The crucial takeaway here is that, due to spontaneous emission alone, the average occupation number of any given mode (including the one that will lase) is very small – significantly less than one photon. The mode is mostly "empty" of spontaneously generated photons at any instant.

Page 18:

Now, let's put this very small number of spontaneous photons per mode into perspective:

- **Compare this with approximately** 10^7 (ten million) photons per mode under 1 mW lasing conditions.

When a laser is operating, say a 1 mW HeNe laser, the number of *stimulated* photons occupying the lasing mode is enormous – on the order of 10^7! This is calculated from the power

$$P = \frac{n_{\text{photons}} \cdot h\nu}{\tau_{\text{cavity}}}$$

where n_{photons} is the number of photons in the cavity, $h\nu$ is photon energy, and τ_{cavity} is photon lifetime. Or, if it's output photons per second, it's $\frac{P}{h\nu}$ photons/sec, and then considering how many are in the mode at once. Regardless of the precise calculation, the number is huge.

So, we have about 0.1 spontaneous photons in the mode versus 10,000,000 stimulated photons in the mode.

This comparison leads to a vital conclusion:

The spontaneous contribution is negligible to the *amplitude* (or intensity) of the laser output once it's lasing. Adding 0.1 random photons to a field of 10,000,000 coherent photons hardly changes the total number or the total power.

However, this tiny spontaneous contribution is *crucial for the phase* of the laser light. Each of those 0.1 "average" spontaneous photons (or more accurately, each spontaneous photon event that adds a photon to the mode) carries a random phase. When it adds to the existing strong, coherent field of stimulated photons, it slightly perturbs the phase of that total field. It's the accumulation of these tiny, random phase kicks from ongoing spontaneous emission that leads to phase diffusion and ultimately defines the fundamental Schawlow-Townes linewidth.

So, spontaneous emission: insignificant for amplitude noise in an operating laser, but paramount for phase noise and the fundamental linewidth.

Page 19:

We now shift our focus to another type of fundamental noise: **Photon-Number (Shot) Noise and its relation to Amplitude Stabilisation.** While we just concluded that spontaneous emission is key for phase noise, let's briefly examine noise in the photon number itself.

* The **Output power,** P, of a laser relates to the average output photon rate, n- bar ($\bar{n}$). The equation is:

$$\bar{n} = \frac{P}{h\nu_{\mathrm{L}}}$$

That is, $\bar{n} = \frac{P}{h\nu_{\mathrm{L}}}$. Here, **n- bar (**$\bar{n}$) is the average number of photons emitted by the laser per unit time (e.g., photons per second). P is the laser's output power (e.g., in Watts or Joules per second). $h\nu_{\mathrm{L}}$ is the energy of a single laser photon (h is Planck's constant, ν_{L} is the laser frequency). This equation simply states that the average photon rate is the total energy per second (power) divided by the energy per photon.

* Let's look at an example: **A** $1\,\mathrm{mW}$ laser where the photon energy $h\nu_{\mathrm{L}}$ is approximately $2\,\mathrm{eV}$. (Note: $2\,\mathrm{eV}$ corresponds to a wavelength of about 620 nanometers, which is in the visible red range, typical for lasers like HeNe). For these parameters, the slide states: **n- bar is approximately** $8 \times 10^{15}\,\mathrm{s}^{-1}$ ($\bar{n} \approx 8 \times 10^{15}\,\mathrm{s}^{-1}$). Let's quickly verify this. $1\,\mathrm{mW} = 10^{-3}\,\mathrm{J/s}$. $1\,\mathrm{eV} \approx 1.602 \times 10^{-19}\,\mathrm{J}$. So, $2\,\mathrm{eV} \approx 3.204 \times 10^{-19}\,\mathrm{J}$.

$$\bar{n} = \frac{10^{-3}\,\mathrm{J/s}}{3.204 \times 10^{-19}\,\mathrm{J/photon}} \approx 3.12 \times 10^{15}\,\mathrm{photons/s}$$

The slide's value of $8 \times 10^{15}\,\mathrm{s}^{-1}$ is about 2.5 times larger. This discrepancy might arise if "P" in the slide's context refers to intracavity power which can be higher than output power, or if n- bar here includes a factor related to cavity Q or lifetime. For our purposes, the exact number isn't as critical as the fact that it's a very large number of photons per second. Let's proceed with the slide's value.

* Now, for the **Statistical distribution of emitted photons in a one- second observation window (for a coherent state, which is a good model for an ideal laser field): it obeys the Poisson law.** A coherent state has photon number statistics that follow a Poisson distribution. This means that if you count the number of photons arriving in a fixed time interval, that number will fluctuate from interval to interval around the mean value $\bar{n}$, according to Poisson statistics.

Page 20:

Continuing with the Poisson statistics of photon emission:

The Poisson probability distribution function, $p(n)$, gives the probability of observing exactly 'n' photons in a given time interval when the average number of photons expected in that interval is $\bar{n}$. The formula is:

$$p(n) = \frac{\bar{n}^n \, e^{-\bar{n}}}{n!}$$

Here, e is Euler's number (the base of natural logarithms), and $n!$ is the factorial of n.

Now, let's consider the fluctuations:

* The **Relative root-mean-square (rms) fluctuation due to this photon statistics (which is shot noise)** is given by:

The square root of (the expectation value of (capital Delta n squared)) divided by n-bar, equals one divided by (the square root of n-bar), which is much, much less than 1.

That is, $\frac{\sqrt{\langle(\Delta n)^2\rangle}}{\bar{n}} = \frac{1}{\sqrt{\bar{n}}} \ll 1$.

Let's break this down:

* $\Delta n = n - \bar{n}$ is the deviation of the actual photon number 'n' from the average '$\bar{n}$'. * $\langle(\Delta n)^2\rangle$ is the variance of the photon number, which for a Poisson distribution is equal to the mean, $\bar{n}$. * So, $\sqrt{\langle(\Delta n)^2\rangle} = \sqrt{\bar{n}}$ is the standard deviation of the photon number. * The relative fluctuation is then $\frac{\sqrt{\bar{n}}}{\bar{n}} = \frac{1}{\sqrt{\bar{n}}}$.

Since $\bar{n}$ (the average number of photons, say, per coherence time or per measurement interval) is typically very large in a laser beam (e.g., 10^7 photons per mode as mentioned earlier, or even larger rates per second), $\sqrt{\bar{n}}$ is also large, making $\frac{1}{\sqrt{\bar{n}}}$ a very small number. For example, if $\bar{n} = 10^6$, then $\frac{1}{\sqrt{\bar{n}}} = 10^{-3}$ or 0.1%. This means that the relative fluctuations in the laser's amplitude due to shot noise are generally very small.

Furthermore, there's a powerful mechanism in lasers that suppresses amplitude fluctuations:

* **Gain saturation:** This is a self-regulating mechanism. **If an instantaneous fluctuation causes** n (the number of photons in the cavity mode) to rise above its steady-state value, the increased stimulated emission rate depletes the population inversion ($N_2 - N_1$) more rapidly. **As the population inversion is depleted, the gain of the laser medium drops. A lower gain means less amplification, so the photon number** n then tends to fall back towards its steady-state value. Conversely, if n drops, the inversion builds up, increasing the gain, which brings n back up. This negative feedback loop ensures that the amplitude (and thus the photon number) is quite stable and self-corrects against perturbations.

Page 21:

Considering the small relative magnitude of shot noise $\frac{1}{\sqrt{\bar{n}}}$ and the strong stabilizing effect of gain saturation on amplitude, we reach an important conclusion regarding noise contributions to the laser linewidth:

* **Result: Amplitude noise contributes very little to the laser linewidth,** $\Delta\nu_L$. The dominant effect is phase noise.

This is a critical point that shapes our entire approach to understanding the fundamental laser linewidth. While there are fluctuations in the amplitude (intensity) of the laser, these are generally small and well-regulated. Moreover, small changes in amplitude do not directly translate to a significant broadening of the spectral line in the same way that phase changes do.

It's the random fluctuations in the *phase* of the laser's electromagnetic field, primarily driven by spontaneous emission events, that cause the frequency to "jitter" or "diffuse" around its central value. This phase diffusion is what directly leads to the spectral broadening we identify as the laser linewidth.

Therefore, when we delve into the Schawlow-Townes theory, our primary focus will be on understanding the origins and consequences of phase noise, as this is the key to the intrinsic linewidth of a laser. The dash at the end of the slide suggests we're moving on from this specific point, having established the dominance of phase noise.

Page 22:

Now, let's explore a **Geometric Picture – using an Amplitude-Phase Representation** to visualize what's happening with the laser field. This provides a very intuitive way to understand phase diffusion.

• We can represent the laser's electric field using a **Complex electric-field phasor.**

Recall that a real, time-varying field $E(t) = A(t)\cos[\omega_0 t + \phi(t)]$ can be thought of as the real part of a complex phasor $\tilde{E}(t) = A(t)e^{i\phi(t)}\, e^{i\omega_0 t}$. If we move into a frame rotating at the carrier frequency ω_0, we can look at the slowly varying complex amplitude (or phasor):

E tilde equals A times e to the power of (i phi) That is, $\tilde{E} = Ae^{i\phi}$.

Here, **E tilde (**$\tilde{E}$) is the complex phasor representing the field. **A** is its instantaneous real amplitude. **φ** is its instantaneous phase. In a polar diagram (complex plane), this phasor is a vector of length A, making an angle ϕ with the real axis.

Now, consider the behavior of A and ϕ in a real laser:

- **The amplitude A is limited to a narrow annulus, capital Delta A (**δA), by the gain saturation mechanism we discussed.

Imagine this phasor in the complex plane. Because gain saturation strongly stabilizes the amplitude, the tip of the phasor $\tilde{E}$ doesn't just wander anywhere. Its length A is constrained to be very close to its average value. So, if you plotted the trajectory of the phasor's tip, it would mostly stay within a thin ring, or annulus, of radial width δA.

- In contrast, **the phase phi (**ϕ) performs a random walk from 0 to 2 pi (and beyond).

While the amplitude A is tightly controlled, there's no such strong restoring force for the phase ϕ. As we'll see, spontaneous emission causes ϕ to take random steps. Over time, these steps accumulate, and the phase angle can wander all over the 0 to 2π range, and indeed, can accumulate to values much larger than 2π (though it's usually considered modulo 2π).

This leads directly to the concept of:

- **Phase diffusion: This is caused by incremental random "kicks" to the phase phi (**ϕ) each time a spontaneous photon enters the cavity mode.

Each spontaneously emitted photon that gets coupled into the lasing mode adds vectorially to the existing strong coherent field. Since the spontaneous photon has a random phase relative to the coherent field, this addition causes a small, random change in the phase (a "kick") of the total field. These kicks are typically small, but they are continuous and random, leading to the diffusive behavior of ϕ.

Page 23:

This page elaborates on a crucial difference between amplitude and phase dynamics:

* **There is no equivalent restoring force for phase.**

This is a key distinction. For amplitude, as we discussed, gain saturation acts like a restoring force: if the amplitude gets too high, gain drops and pulls it back down; if it gets too low, gain increases and pulls it back up. This keeps the amplitude confined.

However, for the phase ϕ, there isn't such a direct, strong restoring mechanism in a simple laser. If a spontaneous emission event kicks the phase by a small amount, there's nothing inherent in the basic laser process that tries to pull the phase back to its original value. The phase is "free to wander."

Therefore, the phase undergoes an unbounded Brownian motion.

This is a direct consequence of the lack of a restoring force and the presence of random phase kicks. Brownian motion, or a random walk, describes a process where successive steps are random. Over time, the displacement from the origin (in this case, the change in phase from its initial value) can grow without limit.

Mathematically, this is characterized by the **variance of the phase change:**

The variance, given by the expectation value of (the quantity $\phi(t)$ minus $\phi(0)$, squared), is proportional to t.

That is,

$$\langle [\phi(t) - \phi(0)]^2 \rangle \propto t.$$

Here, $\phi(0)$ is the phase at some initial time $t = 0$, and $\phi(t)$ is the phase at a later time t. The mean square change in phase grows linearly with time. This is a hallmark of a diffusive process. The longer you wait, the further the phase is likely to have drifted from its starting point, on average.

This unbounded phase diffusion is precisely what leads to a finite linewidth in the laser's spectrum. If the phase were perfectly stable, the linewidth would be zero (a delta function in frequency). But because the phase wanders, the instantaneous frequency (related to $\frac{d\phi}{dt}$) also fluctuates, leading to a spread of frequencies in the output.

Page 24:

This slide presents a beautiful visual illustration titled **Amplitude-Phase Representation: Phase Diffusion.** This diagram helps to solidify the concepts we've just discussed.

We are looking at a **polar plot in the complex plane.**

The origin is at the center. The horizontal axis can be thought of as the **Real part of the complex field** E ($\mathrm{Re}(E)$), and the vertical axis as the **Imaginary part of** E ($\mathrm{Im}(E)$).

A **large, light gray annulus (a ring)** is drawn, centered at the origin. This annulus represents the region where the tip of the electric field phasor is primarily confined. The radial thickness of this annulus is labeled **delta A (**δA), representing the small fluctuations in the amplitude A allowed by gain saturation. The phasor's length (amplitude A) stays mostly within this ring.

An example **phasor** is drawn as a black arrow originating from the center. Its length is labeled A, and its angle with the positive real axis is labeled **phi (**ϕ). This is our $\tilde{E} = Ae^{i\phi}$.

The most important feature is the **"Blue path"**. This is a jagged, wiggly line that shows the trajectory of the *tip* of the phasor over time. Notice several key things about this path:

– It largely stays **within the gray annulus**, meaning the amplitude A does not change much. – However, the path **wanders significantly in the angular direction (phase** ϕ). It performs a random walk around the circle. – The little blue arrows along the path indicate the random "kicks" or steps the phasor takes. Each step has a random component both radially (affecting A slightly within δA) and tangentially (affecting ϕ significantly).

The caption reads: **"Blue path: Random walk of phasor tip over time."**

This diagram perfectly encapsulates the idea:

1. Amplitude A is relatively stable, confined to a narrow range δA due to gain saturation. 2. Phase ϕ is *not* stable; it diffuses randomly over time due to the accumulation of small, random phase kicks (primarily from spontaneous emission).

This random walk of the phase is the very essence of phase diffusion, and it's the direct cause of the fundamental laser linewidth.

Page 25:

Now we move to formalize this phase diffusion in a **Phase Diffusion Model**, with the goal of **Deriving the Lorentzian Spectrum** that characterizes the laser line.

* First, we **Assume a small time increment, delta t (**δt). Over this small time step, the change in phase is also small. We can write the phase at time $t + \delta t$ as:

$\phi(t + \delta t) = \phi(t) + \delta\phi$ That is, $\phi(t + \delta t) = \phi(t) + \delta\phi$. Here, $\delta\phi$ is the small, random change (or "kick") in phase that occurs during the time interval δt.

* Next, we characterize this random phase increment $\delta\phi$: **delta phi (**$\delta\phi$) is assumed to be a Gaussian random variable with zero mean and variance equal to $2\,D\delta t$. Let's unpack this:

* **Zero mean (**$\langle\delta\phi\rangle = 0$): This means that, on average, a phase kick is equally likely to be positive or negative. There's no preferred direction for the phase change in any individual step.

* **Gaussian:** The probability distribution of these small phase kicks is assumed to be Gaussian. This is often justified by the central limit theorem if $\delta\phi$ results from many even smaller underlying random processes.

* **Variance** $\langle(\delta\phi)^2\rangle = 2\,D\delta t$: This is the crucial part. The variance (mean square value) of the phase kick is proportional to the time interval δt. The proportionality constant is $2\,D$. * D here is a very important parameter called the **Phase-diffusion coefficient**. It quantifies how rapidly the phase diffuses. A larger D means more rapid phase diffusion and larger phase kicks per unit time.

* This leads to the formal definition of the **Phase-diffusion coefficient, D, which has units of Hertz (Hz) if phase is in radians:** The phase-diffusion coefficient D is defined from the long-term behavior of the phase variance that we saw earlier ($\langle[\phi(t) - \phi(0)]^2\rangle \propto t$). Specifically:

D equals one-half times the time derivative of (the expectation value of (the quantity phi of t minus phi of zero, squared)) That is, $D = \frac{1}{2}\frac{d}{dt}\langle[\phi(t) - \phi(0)]^2\rangle$.

Since $\langle[\phi(t) - \phi(0)]^2\rangle = 2\,Dt$ (from the properties of $\delta\phi$ for many steps), taking $\frac{d}{dt}$ gives $2\,D$, and then $\frac{1}{2} \times 2\,D = D$. So this is consistent.

If D has units of Hertz, it implies that the variance $2\,Dt$ has units of radians squared (since phase is in radians). So D itself would have units of $\text{rad}^{2/\text}{s}$. However, D is often quoted directly in Hz when relating it to spectral linewidths. This convention arises because the spectral linewidth (FWHM) will turn out to be $2D$ (if D is HWHM in angular frequency) or $D/(2\pi)$ if D is related to phase variance in $\text{rad}^{2/\text}{s}$ and linewidth in Hz. Let's assume the units of D given as Hz on the slide are consistent with its use in defining the spectral FWHM later.

Page 26:

Now that we have established the model for phase diffusion, characterized by the coefficient D, let's see its consequences for the laser's properties.

* First, the **Time-domain correlation function of the field.**

The phase diffusion process causes the laser's electric field to lose correlation with itself over time. The auto-correlation function, \(G(t) = \langle E^*(0)E(t) \rangle\), describes how similar the field at time* t is to the field at time 0 (where \(E^\) is the complex conjugate). For a field undergoing pure phase diffusion (assuming constant amplitude A), this correlation function decays exponentially:

The expectation value of $E^*(0)E(t)$ equals A^2, times $e^{-D|t|}$.

That is,

$$\langle E^*(0)E(t)\rangle = A^2 e^{-D|t|}.$$

Here, A is the constant amplitude of the field (or E_0 from earlier). The term $e^{-D|t|}$ shows that the correlation decays exponentially with a rate D. The absolute value of t, $|t|$, indicates that the decay is symmetric for positive and negative time lags. The coefficient D here directly determines the rate of this decorrelation. A larger D means faster decorrelation.

Note: For this correlation function to lead to a power spectrum, A should be more like E_0.

* Next, the **Fourier transform of this exponential decay gives the Lorentzian power spectrum.**

This is a standard result from Fourier analysis, often encapsulated in the Wiener-Khinchin theorem, which states that the power spectral density of a stationary random process is the Fourier transform of its autocorrelation function. When you Fourier transform an exponential decay like $e^{-D|t|}$, you get a Lorentzian lineshape in the frequency domain. The power spectrum, $|E(\nu)|^2$, is given as:

The absolute value of $E(\nu)$ squared, equals $\frac{E_0^2 D^2}{(\nu-\nu_0)^2+D^2}$.

That is,

$$|E(\nu)|^2 = \frac{E_0^2 D^2}{(\nu - \nu_0)^2 + D^2}.$$

Let's analyze this Lorentzian function:

* E_0^2 is related to the peak intensity or total power. * ν_0 is the center frequency of the laser line. * D (the phase diffusion coefficient from the exponent of the correlation function) appears in the denominator as D^2. This D represents the **Half Width at Half Maximum (HWHM)** of this Lorentzian spectral peak, in units of frequency (e.g., Hertz). *The numerator* $E_0^2 D^2$ ensures the correct peak height. At resonance ($\nu = \nu_0$), the denominator is D^2, so $|E(\nu_0)|^2 = E_0^2$. Thus, E_0^2 is the peak spectral density. The formula should probably read (Peak Value D^2) / $[(\nu - \nu_0)^2 + D^2]$ or if E_0 is field amplitude, then the total power P would be proportional to E_0^2, and the peak spectral density is $\frac{P}{\pi D}$ or similar. A more standard form is $S(\nu) = \frac{S_{\text{peak}} D^2}{(\nu-\nu_0)^2+D^2}$. So E_0^2 in the numerator of the slide's equation represents the peak spectral density.

* From this Lorentzian lineshape, the **Full Width at Half Maximum (FWHM) of the spectrum equals:**

capital Delta nu sub L equals 2D

That is,

$$\Delta\nu_{\mathrm{L}} = 2\,D.$$

This is a fundamental result: if D is the HWHM of the Lorentzian (which it is, as it appears squared with the frequency term in the denominator), then the full width at half maximum is twice that value.

So, the laser linewidth $\Delta\nu_{\mathrm{L}}$ is directly proportional to the phase diffusion coefficient D. Understanding and calculating D is therefore key to determining the fundamental laser linewidth.

Page 27:

This page offers an important clarification regarding terminology, which can often be a source of confusion:

By historical convention, many texts quote half-width at half-maximum, D; we shall use full width, capital $\Delta\nu_{\mathrm{L}}$ ($\Delta\nu_{\mathrm{L}}$).

This is a crucial point of attention.

When discussing spectral lineshapes, especially Lorentzians:

* The parameter D, as it appeared in the denominator of our Lorentzian power spectrum ($(\nu - \nu_0)^2 + D^2$) and as the decay constant in the time-domain correlation function $e^{-D|t|}$, represents the **Half Width at Half Maximum (HWHM)** of the spectral line, assuming D is in frequency units (like Hertz).

* The **Full Width at Half Maximum (FWHM)**, which is the more common experimental measure of linewidth and is denoted here as $\Delta\nu_{\mathrm{L}}$, is therefore **twice the HWHM**. So, $\Delta\nu_{\mathrm{L}} = 2\,D$.

The slide explicitly states that while some literature might refer to "D" as "the linewidth" (implying HWHM), this course or this particular discussion will consistently use $\Delta\nu_{\mathrm{L}}$ to mean the full width. This is good practice for clarity. So, whenever we derive or discuss D, remember that to get the actual, full laser linewidth $\Delta\nu_L$, we will multiply D by two.

The triple dash at the end suggests this is just a point of clarification before we proceed.

Page 28:

Now we delve into the **Physical Origin of the Diffusion Coefficient D.** We've defined D and seen its role in the lineshape, but where does it come from fundamentally?

* The first bullet point gives us the microscopic starting point: Each spontaneously emitted photon that enters the lasing mode adds an uncertain phase, $\delta\phi$, which is on the order of $\frac{1}{\sqrt{\bar{n}}}$. Here, $\bar{n}$ is the average number of coherent photons already present in the lasing mode. When a single spontaneous photon (with random phase and an amplitude roughly corresponding to "one photon") adds to the strong existing field (containing $\bar{n}$ photons), the resulting perturbation to the *phase* of the total field is inversely proportional to the amplitude of that existing field. The field amplitude is proportional to $\sqrt{\bar{n}}$. Thus, the phase kick $\delta\phi$ scales as $\frac{1}{\sqrt{\bar{n}}}$. This means the more photons already in the mode, the smaller the phase kick from a single spontaneous event – the coherent field is "stiffer" against phase perturbations if it's stronger.

* Next, how often do these phase-perturbing events occur? **The Rate of such events (i.e., the rate of these phase kicks) is proportional to the spontaneous photon rate into the mode, which we can call** N_{sp}. N_{sp} here is the number of spontaneously emitted photons that are successfully coupled into the specific lasing mode per unit time (units of s^{-1}). The more frequently these spontaneous photons arrive, the more frequently the phase is kicked.

* Now, putting it together and **Including the cavity photon lifetime,** τ_c, and considerations of output coupling, yields a proportionality for D: The slide states: **D is proportional to** N_{sp}, times $\frac{1}{2\pi\bar{n}\tau_c}$. That is,

$$D \propto \frac{N_{sp}}{2\pi\bar{n}\tau_c}.$$

Let's examine the terms and their roles: * N_{sp} (spontaneous photon rate into the mode, s^{-1}): As N_{sp} increases, D increases. More spontaneous photons mean more phase kicks, faster diffusion. This makes sense. * $\bar{n}$ (average number of coherent photons in the mode): As $\bar{n}$ increases, D decreases (since $\bar{n}$ is in the denominator). A stronger coherent field is more resistant to phase perturbations from individual spontaneous photons. This also makes sense. * τ_c (cavity photon lifetime, s): As τ_c increases (meaning a higher Q cavity, narrower cavity resonance), D decreases. A longer photon lifetime means photons stay in the cavity longer, and the field has more "memory." A narrow cavity resonance (long τ_c) also means it's "harder" for off-resonant spontaneous photons to be accepted into the mode, or they have less impact. The factor $\frac{1}{2\pi\tau_c}$ is essentially the cavity resonance half-width, $\Delta\nu_c$ (HWHM). So, D is proportional to $\frac{N_{sp}}{\bar{n}}$, scaled by something related to the cavity bandwidth. A common form for the HWHM linewidth D (which is $\Delta\nu_L/2$) is

$$D \approx \frac{N_{sp}}{4\pi\bar{n}} K,$$

where K is an excess noise factor often near 1 for simple cases, and N_{sp} here is the rate. This simplified form doesn't explicitly show τ_c. The Schawlow-Townes formula, which we are building towards, will incorporate these factors more precisely. The proportionality shown on the slide is a step in that direction, highlighting the key physical dependencies. We should be careful about the precise prefactors and dimensional consistency, which will become clearer in the full Schawlow-Townes derivation. The crucial insight here is that D increases with the rate of "bad" spontaneous photons and decreases with the number of "good" coherent photons and with longer cavity lifetimes (narrower cavity resonances).

Page 29:

Let's explore the implications of the dependencies we just discussed for the diffusion coefficient D, and consequently, for the laser linewidth.

* The first bullet point considers the effect of laser power:

Larger output power, P_L (PL), implies a larger average number of photons in the cavity, $\bar{n}$. (Since power is proportional to $\bar{n}/\tau_c$).

A larger $\bar{n}$, as we saw, leads to a smaller diffusion coefficient D (because $\bar{n}$ is in the denominator of the expression for D).

And a smaller D means a narrower linewidth $\Delta\nu_L = 2\,D$.

So, the conclusion is: **Higher laser output power generally leads to a narrower fundamental linewidth.** This is a very important practical consequence. If you want a spectrally pure laser (small $\Delta\nu_L$), operating it at higher power (assuming other factors are optimized) is beneficial, up to other limiting effects not considered here.

* The second bullet point considers the effect of the laser cavity's properties:

A smaller cavity bandwidth, $\Delta\nu_c$ (which corresponds to a higher cavity finesse, F), decreases the fraction of spontaneous photons accepted into the lasing mode.

Think of the cavity as a filter. Spontaneous emission occurs over a broader fluorescence bandwidth $\Delta\nu_D$. The cavity resonance, with width $\Delta\nu_c$, only "accepts" or strongly interacts with those spontaneously emitted photons whose frequencies fall within this narrow $\Delta\nu_c$. If $\Delta\nu_c$ is made smaller (e.g., by using higher reflectivity mirrors, leading to higher finesse and longer photon lifetime τ_c), then a smaller fraction of the total spontaneous emission spectrum will effectively couple into the lasing mode and perturb its phase.

Hence, a smaller $\Delta\nu_c$ also reduces D.

This also makes intuitive sense: a "sharper" cavity filter leads to less phase noise from spontaneous emission, thus a smaller D and a narrower laser linewidth.

These two points – higher power and narrower cavity bandwidth – are key design principles for achieving lasers with very narrow Schawlow-Townes limited linewidths. The triple dash indicates we're moving to define these cavity properties more formally.

Page 30:

Now we turn to **Slide 10: Cavity Properties – Definitions and Relations.** Understanding these is essential because, as we just saw, the cavity plays a crucial role in determining the laser linewidth.

* First, let's define the **Resonator half-width, capital Delta nu sub c (**$\Delta\nu_\mathrm{c}$). This $\Delta\nu_\mathrm{c}$ is a measure of the sharpness of the cavity's resonance peaks when you plot its transmission or stored energy as a function of frequency. It's **determined by mirror losses (absorption and scattering) and output coupling** (the intentional transmission of light out of the cavity to form the laser beam). The relationship between $\Delta\nu_\mathrm{c}$ and the photon lifetime in the cavity, τ_c, is:

$$\Delta\nu_\mathrm{c} = \frac{1}{2\pi\tau_\mathrm{c}}$$

That is, $\Delta\nu_\mathrm{c} = \frac{1}{2\pi\tau_\mathrm{c}}$. Here, τ_c is the photon lifetime in the cavity (in seconds). It represents the average time a photon, once inside the cavity, will survive before being lost due to mirror imperfections or output coupling. $\Delta\nu_\mathrm{c}$ as defined here is the **Half Width at Half Maximum (HWHM)** of the cavity resonance, in Hertz. A longer photon lifetime τ_c corresponds to a narrower cavity resonance $\Delta\nu_c$, meaning the cavity is more selective in frequency.

* Next, a very important parameter characterizing a resonator is its **Quality factor,** Q. The Q factor is a dimensionless quantity that represents the "quality" of a resonator. A higher Q means lower losses and a sharper resonance.

Page 31:

Continuing with cavity properties:

The **Quality factor, Q**, is defined as:

Q equals nu sub L divided by (2 times capital Delta nu sub c)

That is,

$$Q = \frac{\nu_L}{2\Delta\nu_c}$$

Here: * ν_L is the laser's operating frequency (which is also the center frequency of the cavity resonance we are considering). * $\Delta\nu_c$ is the **Half Width at Half Maximum (HWHM)** of the cavity resonance, as defined on the previous page ($\Delta\nu_c = \frac{1}{2\pi\tau_c}$). * Therefore, $2\Delta\nu_c$ is the **Full Width at Half Maximum (FWHM)** of the cavity resonance.

So, Q is the ratio of the resonant frequency to the FWHM of the resonance. A high Q factor means the FWHM ($2\Delta\nu_c$) is very small compared to the resonant frequency ν_L, indicating a very sharp, selective resonance.

Another important parameter for optical cavities, especially Fabry-Pérot type cavities, is the **Finesse, capital F (F).**

* The Finesse is defined as the ratio of the Free Spectral Range (FSR) of the cavity to the FWHM of its resonance. The FSR is the frequency spacing between adjacent longitudinal modes of the cavity. The slide gives the formula:

Capital F equals (Free Spectral Range) divided by capital Delta nu sub c, which equals c divided by (2 times d times capital Delta nu sub c).

That is,

$$F = \frac{\text{FSR}}{\Delta\nu_\text{c}} = \frac{c}{2\,d\,\Delta\nu_\text{c}}$$

Let's be careful here with $\Delta\nu_\text{c}$. If $\Delta\nu_\text{c}$ from the Q-factor definition (and page 30) is HWHM, then the FWHM of the cavity resonance is $2\Delta\nu_\text{c}$. The Free Spectral Range (FSR) for a simple two-mirror cavity of length d is $\frac{c}{2\,d}$. So, Finesse $F = \frac{\text{FSR}}{\text{FWHM}_\text{cavity}} = \frac{c/(2\,d)}{2\Delta\nu_{c,\text{HWHM}}}$. This would be

$$F = \frac{c}{4\,d\,\Delta\nu_{c,\text{HWHM}}}$$

However, the slide's formula $F = \frac{c}{2\,d\,\Delta\nu_\text{c}}$ implies that the $\Delta\nu_\text{c}$ used *in this specific finesse formula* is the **Full Width at Half Maximum (FWHM)** of the cavity resonance, not the HWHM. This is a common convention for finesse definition. So, if $\Delta\nu_{c,\text{FWHM}}$ is the full width of the cavity resonance, then $F = \frac{\text{FSR}}{\Delta\nu_{c,\text{FWHM}}}$. Let's assume for this finesse formula on the slide, $\Delta\nu_\text{c}$ represents FWHM. It's important to be aware of whether $\Delta\nu_\text{c}$ refers to HWHM or FWHM as conventions can vary. For the Schawlow-Townes derivation, $\Delta\nu_\text{c}$ often refers to the FWHM of the *passive* cavity.

Now, the significance of these parameters for laser linewidth: * **A high Q or a high Finesse (F) implies a narrow cavity resonance width** $\Delta\nu_\text{c}$ (FWHM). **This, in turn, filters out most of the broad spectrum of spontaneously emitted photons**, only allowing those very close to the cavity resonance to build up. This filtering action is **key to achieving a small laser linewidth, capital Delta nu sub L (**$\Delta\nu_\text{L}$). Essentially, a high-Q, high-Finesse cavity is more "choosy" about which spontaneous photons it allows to interact strongly with the mode, thus reducing the phase noise.

Page 32:

This page shows a diagram illustrating a simple Fabry-Pérot cavity, which is the basis for many lasers. The slide indicates "[IMAGE REQUIRED: Simple Fabry-Pérot cavity drawing with length d, reflectivities R_1, R_2, indicating τ_c, photon paths, and output coupling.]" I will describe the provided image.

The diagram depicts a basic optical resonator.

- We see two parallel mirrors, represented as gray rectangular blocks.
- The left mirror is labeled **R sub 1 (**R_1), indicating its reflectivity.
- The right mirror is labeled **R sub 2 (**R_2), indicating its reflectivity.

- The distance between the two mirrors is labeled d, representing the length of the cavity.

- Inside the cavity, between the mirrors, the label **tau sub c (**τ_c) (Photon Lifetime) is present, signifying that this cavity structure is characterized by a certain average time a photon will exist within it before being lost or coupled out.

- A **red arrow (light ray)** is shown incident on the left mirror (R_1) from outside the cavity.

- A portion of this light enters the cavity and is shown as a red line **bouncing back and forth multiple times** between R_1 and R_2. This represents the resonant light path within the cavity.

- Each time the light hits the right mirror (R_2), a fraction of it is transmitted through. This transmitted light is shown as a red arrow exiting the cavity to the right, labeled **Output Coupling**. This is the useful laser beam.

This simple diagram visually captures the essential elements of a Fabry-Pérot resonator: two mirrors forming a cavity of length d. The reflectivities R_1 and R_2 (along with any other losses) determine the photon lifetime τ_c, and consequently, the cavity's Q-factor, finesse (F), and resonance width

($\mathit{\Delta}\nu_{\mathrm{c}}$). The output coupling through one of the mirrors (typically R_2) allows the laser light to be extracted. The Schawlow-Townes linewidth is fundamentally tied to the properties of such a cavity, interacting with the gain medium placed within it.

Page 33:

We are now ready to begin to begin the **Step-by-Step Derivation of the Schawlow-Townes Limit**, which is a landmark result in laser physics. This is presented on Slide 11.

The derivation **Starts from the phase-diffusion width (which is the half-width)** D.

Recall that $\mathit{\Delta}\nu_{\mathrm{L}}$ (the full linewidth) = $2\,D$. So, D is the HWHM of the laser's spectral line.

A key formula, incorporating the factors we've discussed, for this phase-diffusion half-width D is given as:

D equals (1 divided by 4 pi) times (h nu sub L divided by P sub L) times (capital Delta nu sub c, squared) times (the quantity N sub s p plus N sub t h plus 1).

That is,

$$D = \frac{1}{4\pi} \cdot \frac{h\nu_{\mathrm{L}}}{P_{\mathrm{L}}} \cdot (\mathit{\Delta}\nu_{\mathrm{c}})^2 \cdot \left(N_{sp} + N_{th} + 1\right)$$

Let's break down the terms in this crucial starting equation for D (the laser line HWHM):

D: The half-width at half-maximum of the laser emission spectrum, in Hertz.

h: Planck's constant.

ν_{L}: The laser frequency, in Hertz. So, $h\nu_{\mathrm{L}}$ is the energy of a single laser photon.

P_L: The output power of the laser, in Watts. The ratio $\frac{h\nu_\text{L}}{P_\text{L}}$ has units of (Joules) / (Joules/second) = seconds. It represents the energy per photon divided by the energy per second, effectively an inverse photon rate scaled by h. More accurately, $\frac{P_\text{L}}{h\nu_\text{L}}$ is the photon output rate. So $\frac{h\nu_\text{L}}{P_\text{L}}$ is 1 over this rate.

$\Delta\nu_\text{c}$: This is the **Full Width at Half Maximum (FWHM)** of the *passive optical resonator's* resonance, in Hertz. (Note: some derivations use HWHM for $\Delta\nu_\text{c}$ and adjust prefactors. We must be consistent with how $\Delta\nu_\text{c}$ relates to Q or τ_c.) The $(\Delta\nu_\text{c})^2$ term shows a strong dependence on the cavity quality.

$\left(N_{sp} + N_{th} + 1\right)$: This is a dimensionless factor representing the effective number of noise photons. We'll define these terms on the next page.

○ N_{sp} relates to spontaneous emission. ○ N_{th} relates to thermal (blackbody) radiation within the cavity mode. ○ The **"+1"** is profoundly important – it represents the contribution of vacuum fluctuations or zero-point energy of the electromagnetic field. It ensures there's a fundamental linewidth even if N_{sp} and N_{th} were zero (which they aren't). This term is sometimes called the "one noise photon per mode" contribution that is always present.

The term $\frac{h\nu_\text{L}}{P_\text{L}}$ $(\Delta\nu_\text{c})^2$ has units of (seconds) (Hertz squared) = Hertz. So D is in Hertz, which is correct for a linewidth (HWHM).

The factor $\frac{1}{4\pi}$ is a numerical prefactor that arises from the detailed derivation relating phase variance to spectral properties.

The "where" at the bottom of the slide indicates that the next page will clarify N_{sp} and N_{th}.

Page 34:

Continuing the derivation of the Schawlow-Townes limit, let's define the terms in the noise factor $\left(N_{\text{sp}} + N_{\text{th}} + 1\right)$ from the equation for D on the previous page.

* N_{sp} represents the number of spontaneous photons per second (s^{-1}) in the lasing mode. More precisely, N_{sp} in this context is often taken as

$$n_{\text{sp}} = \frac{N_2}{N_2 - \frac{N_1 g_2}{g_1}}$$

which is the "spontaneous emission factor" or "population inversion factor." It's a dimensionless quantity, typically greater than or equal to 1, that accounts for the fact that not all atoms are in the upper state (incomplete inversion). If N_{sp} were truly a rate (s^{-1}), the units in the D formula would be incorrect as discussed. So, let's interpret N_{sp} here as this dimensionless factor related to the degree of inversion and how effectively spontaneous emission contributes noise relative to the stimulated emission process. For an ideal four-level laser, N_{sp} can approach 1. For a three-level laser, it can be larger.

* N_{th} represents the number of thermal photons per second (s^{-1}) in the lasing mode. These are photons from blackbody radiation corresponding to the temperature of the cavity. The slide correctly notes that N_{th} is **(negligible at optical frequencies, where** N_{th} is much, much less than 1).
The average number of thermal photons per mode is given by the Planck distribution Bose-Einstein factor,

$$\frac{1}{e^{\frac{h\nu_{\text{L}}}{kT}} - 1}$$

At optical frequencies ($h\nu_{\text{L}} \gg kT$, where k is Boltzmann's constant and T is temperature), this factor is extremely small. So, for most lasers operating in the visible or near-infrared, the contribution of thermal photons to the linewidth is negligible compared to spontaneous emission and vacuum fluctuations. We can usually set $N_{\text{th}} \approx 0$. Again, if N_{th} were a rate, unit issues arise. So, similar to N_{sp}, N_{th} here should be understood as the effective number of thermal noise photons, which is dimensionless and very small.

* The **"+1" term in** $\left(N_{\text{sp}} + N_{\text{th}} + 1\right)$ counts vacuum fluctuations. This is the contribution from the zero-point energy of the electromagnetic field in the lasing mode. It's always present, even at zero temperature and with perfect inversion. This "1" ensures that there's a fundamental quantum limit to the linewidth that cannot be surpassed.

Now, we take the next step:

* **Multiplying** D (the half-width) by 2 converts it to the full linewidth, capital $\varDelta\nu_L$ ($\varDelta\nu_L$). $\varDelta\nu_L = 2\,D$. Using the expression for D from the previous page:

$$D = \frac{1}{4\pi}\frac{h\nu_{\text{L}}}{P_{\text{L}}}(\varDelta\nu_{\text{c}})^2\left(N_{\text{sp}} + N_{\text{th}} + 1\right)$$

So,

$$\varDelta\nu_L = 2 \times \frac{1}{4\pi}\frac{h\nu_{\text{L}}}{P_{\text{L}}}(\varDelta\nu_{\text{c}})^2\left(N_{\text{sp}} + N_{\text{th}} + 1\right)$$

$$\varDelta\nu_L = \frac{1}{2\pi}\frac{h\nu_{\text{L}}}{P_{\text{L}}}(\varDelta\nu_{\text{c}})^2\left(N_{\text{sp}} + N_{\text{th}} + 1\right)$$

* However, the formula presented on the slide for $\varDelta\nu_L$ is:

$$\varDelta\nu_L = 2\,D = \frac{\pi h\nu_{\text{L}}(\varDelta\nu_{\text{c}})^2\left(N_{\text{sp}} + N_{\text{th}} + 1\right)}{2P_{\text{L}}}$$

That is,

$$\varDelta\nu_L = 2\,D = \frac{\pi h\nu_{\text{L}}(\varDelta\nu_{\text{c}})^2\left(N_{\text{sp}} + N_{\text{th}} + 1\right)}{2P_{\text{L}}}$$

Let's compare the prefactors. My direct doubling of D gives a prefactor of $\frac{1}{2\pi}$. The slide's formula for $\Delta\nu_L$ has a prefactor of $\frac{\pi}{2}$. The ratio is

$$\frac{\pi/2}{1/(2\pi)} = \pi^2$$

This means there's a factor of π^2 difference between simply doubling the D formula from the previous slide and the $\Delta\nu_L$ formula presented here. This often arises from choices in defining $\Delta\nu_c$ (e.g., angular vs. regular frequency, or whether it represents the cold cavity or includes effects of the gain medium like dispersion). The Schawlow-Townes original paper and various textbook derivations can have slightly different forms depending on these definitions. For consistency, we will proceed with the formula for $\Delta\nu_L$ as written on *this* slide, assuming it is the intended expression for this derivation path. The key dependencies (on $h\nu_L$, P_L, $(\Delta\nu_c)^2$, and the noise photon sum) are the most critical aspects.

Page 35:

Now we approach the celebrated result.

The **Minimum possible value for the laser linewidth,** $\Delta\nu_L$, is reached under ideal conditions. These conditions are:

When exactly one spontaneous photon per cavity lifetime effectively initiates the stimulated chain that leads to a phase perturbation. This is encapsulated by setting the term $(N_{sp} + N_{th} + 1)$ to its minimum effective value.

The slide states this minimum is achieved when $N_{sp} = 1$ (if N_{sp} is the population factor, its minimum is 1 for a perfect 4-level system), and $N_{th} \approx 0$ (which is true for optical frequencies).

If we take the expression for $\Delta\nu_{\mathrm{L}}$ from the previous page:

$$\Delta\nu_{\mathrm{L}} = \frac{\pi h \nu_{\mathrm{L}} (\Delta\nu_{\mathrm{c}})^2 \left(N_{sp} + N_{th} + 1\right)}{2P_{\mathrm{L}}}$$

And if the term $\left(N_{sp} + N_{th} + 1\right)$ effectively becomes **2** in the most fundamental quantum limit (where the "+1" for vacuum fluctuations is always there, and the N_{sp} term effectively contributes another "1" for the minimal spontaneous emission noise in an ideal laser), then the formula simplifies. This "factor of 2" for the total effective noise photons (one "real" spontaneous photon + one "vacuum" photon) is a common feature in refined Schawlow-Townes derivations.

This simplification **yields the celebrated Schawlow-Townes formula** for the fundamental quantum-limited laser linewidth, denoted here as $\Delta\nu_{ST}$:

The formula is presented in a box: **capital Delta nu sub S T equals (pi h nu sub L, times capital Delta nu sub c squared) divided by (P sub L).**

That is,

$$\Delta\nu_{ST} = \frac{\pi h \nu_{\mathrm{L}} (\Delta\nu_{\mathrm{c}})^2}{P_{\mathrm{L}}}$$

Comparing this to the formula for $\Delta\nu_{\mathrm{L}}$ on the previous page, we see that the term $\left(N_{sp} + N_{th} + 1\right)$ has effectively been replaced by 2 to arrive at this final expression for $\Delta\nu_{ST}$.

$$\Delta\nu_{\mathrm{L}} = \frac{\pi h \nu_{\mathrm{L}} (\Delta\nu_{\mathrm{c}})^2}{2P_{\mathrm{L}}} \cdot \left(N_{sp} + N_{th} + 1\right)$$

If $\left(N_{sp} + N_{th} + 1\right) \to 2$, then

$$\Delta\nu_{ST} = \frac{\pi h\nu_{\mathrm{L}}(\Delta\nu_{\mathrm{c}})^2}{2P_{\mathrm{L}}} \cdot 2 = \frac{\pi h\nu_{\mathrm{L}}(\Delta\nu_{\mathrm{c}})^2}{P_{\mathrm{L}}}$$

This is consistent. This is the Schawlow-Townes limit, representing the narrowest possible linewidth a laser can achieve, dictated by quantum noise.

Let's highlight the **Key insight** from this formula:

The fundamental laser linewidth $\Delta\nu_{ST}$ is proportional to the square of the passive cavity linewidth $(\Delta\nu_{\mathrm{c}})^2$, and inversely proportional to the output power P_{L}.

This tells us exactly what to do to design a laser with an extremely narrow intrinsic linewidth:

1. Make the passive cavity resonance $\Delta\nu_{\mathrm{c}}$ as narrow as possible (i.e., use very high reflectivity mirrors, make a high-Q, high-Finesse cavity). The squared dependence means this is very effective.

2. Operate the laser at as high an output power P_{L} as practically possible (while maintaining other desired characteristics).

This formula is a cornerstone of laser physics.

Page 36

Slide 12

Let's look at some **Practical Calculation Examples** using the Schawlow-Townes formula:

$$\Delta\nu_{ST} = \frac{\pi\, h\, \nu_{\mathrm{L}}\, (\Delta\nu_{\mathrm{c}})^2}{P_{\mathrm{L}}}$$

First, for a **Helium-Neon (HeNe) laser operating at a wavelength of 632.8 nanometers:**

Given parameters:

Laser frequency, $\nu_{\mathrm{L}} = 5 \times 10^{14}$ Hz.

(This corresponds to $\lambda = \frac{c}{\nu_{\mathrm{L}}} = \frac{3\times10^{8}\ \mathrm{m/s}}{5\times10^{14}\ \mathrm{Hz}} = 0.6 \times 10^{-6}\ \mathrm{m} = 600\ \mathrm{nm}$, which is close to 632.8 nm).

Passive cavity FWHM linewidth, $\Delta\nu_{\mathrm{c}} = 1\ \mathrm{MHz} = 1 \times 10^{6}$ Hz.

(This is a typical value for a HeNe laser cavity with reasonable mirrors).

Output power, $P_{\mathrm{L}} = 1\ \mathrm{mW} = 1 \times 10^{-3}$ W.

Plugging these into the Schawlow-Townes formula (with $h \approx 6.626 \times 10^{-34}$ J·s):

$$\Delta\nu_{ST} = \frac{\pi\,(6.626 \times 10^{-34}\ \mathrm{J{\cdot}s})\,(5 \times 10^{14}\ \mathrm{s}^{-1})\,(1 \times 10^{6}\ \mathrm{s}^{-1})^{2}}{1 \times 10^{-3}\ \mathrm{J/s}}$$

$$\Delta\nu_{ST} = \frac{\pi \cdot 6.626 \cdot 5 \cdot 10^{-34+14+12}}{10^{-3}}\ \mathrm{Hz}$$

$$\Delta\nu_{ST} = \frac{103.084 \times 10^{-8}}{10^{-3}}\ \mathrm{Hz}$$

$$\Delta\nu_{ST} = 103.084 \times 10^{-5}\ \mathrm{Hz} \approx 1.03 \times 10^{-3}\ \mathrm{Hz}.$$

The slide gives the result: $\Delta\nu_{\mathrm{L}}$ is approximately 1.0×10^{-3} Hz.

This is 1 milliHertz (mHz). Our calculation agrees perfectly. So, the fundamental quantum limit for a typical HeNe laser is incredibly small, around a milliHertz!

Next, for an **Argon-ion laser operating at a wavelength of 488 nanometers:**

Given parameters:

Laser frequency, $\nu_\mathrm{L} = 6 \times 10^{14}$ Hz.

$$(\lambda = \frac{c}{\nu_\mathrm{L}} = \frac{3\times10^{8}\ \mathrm{m/s}}{6\times10^{14}\ \mathrm{Hz}} = 0.5 \times 10^{-6}\ \mathrm{m} = 500\ \mathrm{nm}, \text{ close to 488 nm}).$$

Passive cavity FWHM linewidth, $\Delta\nu_\mathrm{c} = 3$ MHz $= 3 \times 10^{6}$ Hz.

(Argon ion lasers often have shorter, lossier cavities than HeNes, so a broader $\Delta\nu_\mathrm{c}$).

Output power, $P_\mathrm{L} = 1$ W.

(Argon lasers are capable of much higher powers than HeNes).

Plugging these in:

$$\Delta\nu_{ST} = \frac{\pi\,(6.626 \times 10^{-34}\ \mathrm{J{\cdot}s})\,(6 \times 10^{14}\ \mathrm{s^{-1}})\,(3 \times 10^{6}\ \mathrm{s^{-1}})^2}{1\ \mathrm{J/s}}$$

$$\Delta\nu_{ST} = \frac{\pi \cdot 6.626 \cdot 6 \cdot 9 \cdot 10^{-34+14+12}}{1}\ \mathrm{Hz}$$

$$\Delta\nu_{ST} = \pi \cdot 357.804 \times 10^{-8}\,\text{Hz}$$

$$\Delta\nu_{ST} \approx 1124.07 \times 10^{-8}\,\text{Hz} \approx 1.124 \times 10^{-5}\,\text{Hz}.$$

The slide gives the result: $\Delta\nu_\text{L}$ is approximately 1.1×10^{-5} Hz.

This is 11 microHertz (μHz). Again, our calculation agrees very well. Notice that even though the Argon laser has a broader cavity linewidth $\Delta\nu_\text{c}$ (which tends to increase $\Delta\nu_{ST}$), its much higher output power P_L (which tends to decrease $\Delta\nu_{ST}$) results in an even narrower fundamental linewidth than the HeNe example.

These examples dramatically illustrate just how small the Schawlow-Townes quantum limits are for common lasers.

Page 37:

Following these practical calculations, the crucial takeaway is emphasized:

* **These values (the calculated Schawlow-Townes linewidths of milliHertz for HeNe and tens of microHertz for Argon-ion) are many orders of magnitude narrower than technical noise can presently allow.**

This is the stark reality. While quantum mechanics permits these incredibly sharp spectral lines, the real world of engineering and experimental physics imposes much broader linewidths due to "technical noise." These technical noise sources, such as mechanical vibrations of the laser cavity, temperature fluctuations, electronic noise in power supplies, acoustic disturbances, and instabilities in the gain medium, typically broaden the laser line far beyond its Schawlow-Townes limit.

For the HeNe example, we calculated a $\sim 1\,\text{mHz}$ limit. A real HeNe laser might have a linewidth of many kHz or even MHz, so that's a factor of 10^6 to 10^9 broader!

For the Argon-ion laser, $\sim 11\,\mu\text{Hz}$ limit. A real one might be many MHz broad.

This huge gap between the fundamental quantum limit and typical achieved linewidths highlights two things:

1. The profound challenge experimentalists face in trying to approach the quantum limit. It requires extraordinary efforts in stabilization and noise suppression. 2. The Schawlow-Townes limit serves as an ultimate benchmark – a theoretical "brick wall" that we cannot surpass, no matter how perfect our engineering.

Understanding this gap is essential for appreciating the ongoing research and development in ultra-stable lasers.

Page 38:

This slide provides a compelling visual comparison: **Laser Linewidths: Theoretical Limits vs. Typical Achieved Values.** This is a bar chart that graphically illustrates the gap we just discussed.

* The **vertical axis represents Linewidth, capital** $\Delta\nu$ (Δv), in Hertz (Hz), plotted on a logarithmic scale. The scale spans from 10^{-6} Hz (microHertz) at the bottom to 10^{6} Hz (MegaHertz) at the top, covering twelve orders of magnitude.

* The **horizontal axis lists several common types of lasers:** * **HeNe** (Helium-Neon laser) * **Ar-ion** (Argon-ion laser) * **ECDL** (External Cavity Diode Laser) * **Fiber** (Fiber laser) * **TiSap** (Titanium-Sapphire laser)

* For each laser type, there are two bars: * A **blue bar representing the Schawlow-Townes Limit** (the theoretical minimum linewidth). * An **orange bar representing the Typical Achieved linewidth** in practice.

Let's look at the trends:

* **HeNe:** The blue bar (Schawlow-Townes) is very low, around 10^{-3} Hz (milliHertz), consistent with our calculation. The orange bar (Typical Achieved) is much higher, around 10^3 to 10^4 Hz (kiloHertz). A gap of about 6-7 orders of magnitude.

* **Ar-ion:** The blue bar is even lower, around 10^{-5} Hz (tens of microHertz). The orange bar is very high, perhaps 10^5 to 10^6 Hz (hundreds of kHz to MHz). An even larger gap of 10-11 orders of magnitude.

* **ECDL (External Cavity Diode Laser):** Blue bar around 10^{-3} Hz. Orange bar around 10^5 Hz (hundreds of kHz), though state-of-the-art ECDLs can be much narrower.

* **Fiber Lasers:** Blue bar very low, perhaps 10^{-5} Hz or less. Orange bar around 10^3 Hz (kHz) for typical systems, but specialized fiber lasers can get much narrower, even sub-Hertz.

* **Ti:Sapphire Lasers:** Blue bar very low, around 10^{-4} Hz. Orange bar around 10^5 Hz.

The visual message is striking: for all these laser types, the **orange bars (typical achieved linewidths) are dramatically taller (broader) than the blue bars (theoretical Schawlow-Townes limits).** This underscores the fact that in most practical lasers, the observed linewidth is dominated by technical noise, not by the fundamental quantum noise described by Schawlow and Townes. Closing this gap is a major ongoing effort in laser science and technology.

Page 39:

Slide 13:

Now, let's discuss the **Experimental Reality – Comparison with Achieved Linewidths** in more detail.

* For **Typical good single-mode gas lasers (like HeNe or Argon-ion) or dye lasers, achieved with moderate experimental effort:**

The laser linewidth, $\Delta\nu_{\mathrm{L}}$, is typically in the range of 10^4 to 10^6 Hertz ($10^4 - 10^6$ Hz).

This means linewidths from about 10 kiloHertz to 1 MegaHertz. This corresponds to the "orange bars" we saw on the previous graph for many common systems. While this is already very monochromatic compared to other light sources, it's still far from the milliHertz or microHertz Schawlow-Townes limits. "Moderate effort" implies standard laboratory setups without extreme stabilization measures.

* However, with **State-of-the-art ultra-stabilised systems,** the situation can be dramatically improved. These systems often involve:

* Using reference cavities made from materials with ultra-low thermal expansion (ULE) coefficients, like ULE glass or Zerodur, to minimize length fluctuations due to temperature changes. * Sophisticated vibration isolation platforms to shield the laser and cavity from mechanical disturbances. * Advanced electronic feedback (servo) systems to actively lock the laser frequency to the stable reference cavity or to an atomic transition. In such highly engineered systems, it's possible to achieve: $\Delta\nu_{\mathrm{L}} \leq 1$ Hz. This is a phenomenal achievement! Linewidths at the Hertz level or even sub-Hertz level have been demonstrated for various types of lasers when extreme stabilization techniques are employed. This brings us much closer to the fundamental limits for some systems, though still many orders of magnitude away for others where the S-T limit is in the microHertz range.

This shows that while the Schawlow-Townes limit is a theoretical floor, meticulous engineering can significantly reduce technical noise, pushing laser performance towards that fundamental boundary.

Page 40

So, what causes the **Remaining gap between the measured linewidth (even in state-of-the-art systems) and the Schawlow-Townes (ST) limit?**

The slide lists several culprits, all of which are essentially forms of residual technical noise that haven't been perfectly eliminated:

- ***Residual cavity-length fluctuations, often denoted by fluctuations in the optical path length*** $n \cdot d$ (nd). Even with ULE materials and vibration isolation, there will still be some tiny changes in the effective cavity length due to temperature drifts, mechanical creep, or uncompensated vibrations.

- ***Acoustic and electronic noise in servo loops.*** The very feedback systems used to stabilize the laser can themselves introduce noise. Acoustic noise can vibrate components, and electronic noise in detectors, amplifiers, and actuators within the servo loop can translate into frequency fluctuations of the laser.

- ***Environmental temperature drifts inducing index changes.*** Small changes in the ambient temperature can affect the refractive index n of air (if any part of the beam path is in air) or of intracavity optical components, again shifting the cavity's resonant frequency.

These are the persistent enemies of ultra-narrow linewidths.

Recognizing these challenges, **Research continues on advanced techniques to close this gap further and push lasers closer to their quantum-limited performance.** Some of these techniques include:

- **Spectral hole burning:** In some gain media, it's possible to use one laser to "burn" a very narrow spectral hole in the gain profile. A second laser can then be locked to this sharp feature, effectively narrowing its linewidth.

- **Active cancellation:** This involves more sophisticated feedback and feed-forward schemes to sense and actively cancel out noise sources. For example, measuring vibrations and applying counter-vibrations, or sensing temperature drifts and actively correcting them.

• **Cryogenic cavities:** Operating the reference cavity at very low (cryogenic) temperatures can dramatically reduce thermal expansion effects, Brownian motion of the mirror surfaces (a fundamental thermal noise limit in itself, distinct from Schawlow-Townes), and other temperature-sensitive noise sources. This is a frontier area for achieving extremely high frequency stability and narrow linewidths.

The quest for " quieter " lasers is an ongoing and exciting field of research.

Page 41:

We now arrive at a very important distinction, highlighted on **Slide 14: Linewidth vs. Frequency Stability – Do Not Confuse!** These two terms describe different aspects of a laser's spectral purity, and it's crucial to understand the difference.

* First, **Linewidth, capital** $\Delta\nu_\mathrm{L}$: This is the **intrinsic spectral width of the laser line, primarily due to random phase diffusion.** This is what the Schawlow-Townes limit describes. It refers to how "fuzzy" or "spread out" the laser's frequency is at any given instant, or over very short timescales. It's related to the coherence time of the laser – how long the phase remains predictable. A narrow $\Delta\nu_\mathrm{L}$ means high coherence.

* Second, **Frequency stability:** This refers to the **ability of the *line centre*,** ν_0, of the laser spectrum to remain constant over time. It's about how much the *average* frequency of the laser drifts, wanders, or jitters over longer timescales (seconds, minutes, hours, or even longer). A laser could have a very narrow instantaneous linewidth ($\Delta\nu_\mathrm{L}$ is small) but its center frequency ν_0 could be drifting all over the place. Conversely, a laser could have a relatively broad instantaneous linewidth but its center frequency might be, on average, very stable.

These are two distinct figures of merit for a laser.

* The slide then gives some **Example achievements** in frequency stability, which we'll see on the next page. These achievements often involve locking the laser to a very stable external reference.

Page 42:

Continuing with examples of frequency stability:

We often characterize frequency stability by the fractional frequency fluctuation, $\frac{\delta\nu}{\nu}$, where $\delta\nu$ is the typical fluctuation or drift of the center frequency ν over a specified observation time.

* For **dye lasers locked to molecular references** (e.g., an absorption line in iodine vapor): Achieved stabilities can be $\frac{\delta\nu}{\nu} \leq 10^{-15}$. This is an incredible level of stability – one part in 10^{15}! This means if the laser frequency is, say, 5×10^{14} Hz, the fluctuation $\delta\nu$ is less than 0.5 Hz.

* For **cavity-stabilised HeNe lasers or solid-state lasers** (where the laser is locked to an ultra-stable Fabry-Pérot reference cavity, like the ULE cavities mentioned earlier): Stabilities can be even better, reaching $\frac{\delta\nu}{\nu} \leq 10^{-16}$. This is approaching the stability required for the most advanced optical atomic clocks.

Now, the crucial link (or lack thereof) between linewidth and stability:

***A narrow linewidth is a* necessary but not sufficient* condition for excellent frequency stability.**

You generally need a laser with a reasonably narrow intrinsic linewidth ($\Delta\nu_L$) as a starting point because it's easier to lock a "clean" narrow line to a reference. If the laser line itself is very broad and noisy, it's hard to define its center precisely enough for tight locking.

However, just having a narrow intrinsic linewidth (e.g., a laser that is close to its Schawlow-Townes limit if it were perfectly isolated) does *not* guarantee that its center frequency ν_0 will be stable over time. That center frequency can still drift due to slow changes in the cavity length, temperature, etc., unless active stabilization measures are taken.

Therefore, for excellent long-term frequency stability, active feedback and referencing (to an atomic transition or a stable cavity) are absolutely required. These servo systems work to keep the laser's center frequency locked to the reference, correcting for drifts.

Page 43:

This slide provides an excellent visual: **Distinction Between Linewidth and Frequency Stability**, showing two plots that clearly illustrate the difference.

Both plots show **Intensity on the vertical axis** and **Frequency (**ν) on the horizontal axis.

* **Plot A is titled: "Broad Linewidth, Good Stability."** * This plot shows a **single, wide, bell-shaped curve.** The width of this curve, indicated by a red double arrow and labeled $\Delta\nu_L$(broad), represents a broad instantaneous laser linewidth. So, at any given moment, the laser's emission is spread over a relatively wide range of frequencies. * However, a vertical dashed blue line indicates the center frequency, ν_0. The label **"Good Stability"** (in green) and the fact that only a single, fixed peak is shown imply that this center frequency ν_0 is very stable over time; it doesn't drift. * So, this laser is "fuzzy" but its average color is very constant.

* **Plot B is titled: "Narrow Linewidth, Poor Stability."** *This plot shows **several snapshots (solid blue line, and fainter dotted lines) of a very** **narrow* bell-shaped curve.** The width of this narrow peak, indicated by a red double arrow and labeled $\Delta\nu_L$(narrow), represents a narrow instantaneous laser linewidth. So, at any given moment, the laser is emitting a very pure frequency. * However, the center frequency of this narrow peak, labeled $\nu_0(t)$, is shown to be shifting over time. The purple

double arrow labeled $\delta\nu$(drift) indicates the range over which this center frequency wanders. The label **"Poor Stability"** (in red) highlights this drift.
* So, this laser is instantaneously "sharp" but its average color changes or drifts over time.

This pair of diagrams is extremely effective in communicating that linewidth (the instantaneous "fuzziness") and stability (the long-term constancy of the center frequency) are independent characteristics of a laser. For many high-precision applications, you need *both* a narrow linewidth *and* good frequency stability.

Page 44:

Finally, let's consider the **Implications for Ultra-High-Resolution Spectroscopy** and other applications that stem from understanding and controlling laser linewidths.

This is Slide 15.

* The first point is directly relevant to spectroscopists: **The resolution limit of a spectroscopic measurement is often set by the probe-laser linewidth rather than the intrinsic linewidth of the sample (e.g., an atomic or molecular transition).** If your laser line is broader than the natural linewidth of the transition you are studying, you will not be able to resolve the true shape or width of the sample's spectral feature. The measured spectrum will be a convolution of the sample's spectrum and your laser's spectrum, effectively "smeared out" by the laser. Therefore, to perform ultra-high-resolution spectroscopy, you need a laser whose linewidth $\Delta\nu_{\mathrm{L}}$ is significantly narrower than the features you aim to resolve.

* The connection to coherence length: **A long coherence length, which results from a narrow linewidth, enables powerful interferometric techniques.** We saw that $L_{\mathrm{c}} \approx \frac{c}{\pi \Delta\nu_{\mathrm{L}}}$. For very narrow linewidths, L_{c} can be enormous. **For example, gravitational-wave detectors (like LIGO and Virgo) use lasers with sub-Hertz linewidths.** This gives them coherence

lengths of hundreds of thousands of kilometers, which is essential for maintaining phase coherence of the light over the long arms of their interferometers (several kilometers long) and for detecting the incredibly tiny changes in arm length caused by passing gravitational waves.

* Knowledge of the Schawlow-Townes (ST) limit has practical engineering implications: **Knowledge of the ST limit guides the engineering of laser systems designed for ultra-high precision and stability.** It tells us what to focus on when trying to achieve the narrowest possible linewidths. Specifically, it guides the design of:

Page 45:

Continuing with how knowledge of the Schawlow-Townes limit guides engineering:

* **Cavity finesse and mirror coatings:** The ST limit shows that linewidth is proportional to $(\Delta\nu_{\text{c}})^2$, the square of the passive cavity linewidth. To make $\Delta\nu_{\text{c}}$ small, we need high finesse, which in turn requires high-reflectivity mirror coatings with minimal losses. So, efforts in developing ultra-high quality mirrors are directly motivated by the desire to reduce $\Delta\nu_{\text{c}}$ and approach the ST limit.

* **Output coupling:** This involves a trade-off. The ST formula shows linewidth is inversely proportional to output power P_{L}. To get high P_{L}, you might need significant output coupling. However, higher output coupling generally means lower finesse and broader $\Delta\nu_{\text{c}}$ (as more energy is lost per round trip). So, there's an optimization problem to balance high power (for low ST linewidth) with high finesse (also for low ST linewidth via small $\Delta\nu_{\text{c}}$). Understanding the ST framework helps in making these design choices.

Pump-noise suppression and thermal design: *While the ST limit itself assumes these technical noise sources are absent, striving to reach the ST limit means aggressively tackling them. Suppressing noise in the pump source (which can affect* N_{sp} or cause other fluctuations) and meticulous thermal design of the laser cavity (to minimize $\Delta\nu_c$ fluctuations due to temperature) are critical steps in reducing the actual *linewidth towards the* fundamental* ST limit.

Beyond spectroscopy and interferometry, there are other profound implications:

* **Understanding the phase noise foundation (which leads to the ST limit) is critical for next-generation optical clocks targeting** 10^{-18} (10^{-18}) fractional stability. Optical atomic clocks are the most precise instruments ever built. Their performance relies on lasers with exceptional stability and narrow linewidths to probe ultra-narrow atomic transitions. The fundamental limit to how well these lasers can perform is ultimately tied to phase noise mechanisms like those described by Schawlow and Townes. Reaching stabilities of 1 part in 10^{18} (which means being able to keep time accurately to within a second over the age of the universe) requires pushing laser technology to its absolute quantum limits.

And a final concluding thought:

* **The Schawlow-Townes framework remains a cornerstone for designing, diagnosing, and pushing laser systems toward quantum-limited performance.** Even though the original formula has seen refinements and extensions (e.g., to include bad-cavity effects, semiconductor laser specifics like the linewidth enhancement factor), its core insights about the role of spontaneous emission, cavity properties, and power remain fundamentally important. It provides the language and the conceptual tools for physicists and engineers working to create ever more perfect sources of coherent light.

The triple dash indicates the end of this section on laser linewidths. This has been a deep dive, from basic definitions to fundamental quantum limits and practical implications. I hope this has given you a thorough understanding of why lasers have linewidths and what determines their ultimate spectral purity.

And with that, we conclude our detailed exploration of the linewidths of single-mode lasers, culminating in the elegant and profoundly insightful Schawlow-Townes limit. We've journeyed from the initial, perhaps unsettling, question of why an idealized source like a laser should even possess a non-zero linewidth, all the way to understanding the quantum mechanical origins of that width.

We've seen that the random phase kicks imparted by spontaneously emitted photons, coupled into the lasing mode, lead to phase diffusion – a random walk of the laser field's phase in the complex plane. This diffusion, in turn, manifests as a Lorentzian lineshape in the frequency domain, with a full width at half maximum that is fundamentally limited. The Schawlow-Townes formula provided us with a quantitative prediction for this limit, revealing its dependence on crucial parameters like the laser power, the photon energy, and, very significantly, the square of the passive cavity linewidth.

The numerical examples for HeNe and Argon-ion lasers dramatically illustrated just how incredibly narrow these fundamental limits are – milliHertz down to microHertz! Yet, as the comparison charts and discussion of experimental reality showed, achieving these limits is an ongoing, monumental challenge due to the pervasive nature of technical noise. Vibrations, thermal effects, electronic noise – these are the practical hurdles that often dominate the observed linewidths by many orders of magnitude.

However, the Schawlow-Townes limit is far from being a mere academic curiosity. It serves as an indispensable benchmark, guiding the sophisticated engineering efforts in ultra-stabilization techniques, from ULE cavities and vibration isolation to advanced servo-locking schemes and cryogenic systems. The pursuit of this quantum limit is what drives innovation in areas like ultra-high-resolution spectroscopy, gravitational wave detection, and, perhaps most notably, the development of next-generation optical atomic clocks striving for unprecedented levels of precision.

It's also crucial to carry forward the distinction between the intrinsic linewidth $\Delta\nu_L$ and long-term frequency stability $\delta\nu/\nu$. A narrow linewidth is necessary for good stability, but active referencing and feedback are paramount for ensuring that the laser's center frequency remains anchored over time.

So, the Schawlow-Townes framework isn't just a formula; it's a conceptual cornerstone. It helps us understand the ultimate performance boundaries imposed by quantum mechanics and provides a clear roadmap for designing, diagnosing, and continually pushing laser technology towards that elusive quantum-limited frontier. This understanding is vital for anyone engaged in advanced laser spectroscopy or applications demanding the highest spectral purity and stability.

That brings us to the end of the material presented in these slides on laser linewidths. We'll pause here for this topic. In our next lecture, we will build upon these concepts as we explore further aspects of laser behavior and their spectroscopic applications. Are there any immediate, burning questions on what we've covered today regarding laser linewidths?